CERTIFIED SELECTIVE TRACKING OF MOVING KKT SYSTEMS

Anonymous authors
Paper under double-blind review

ABSTRACT

A small KKT residual can arrive after a deadline or correspond to the wrong active set or stationary branch. We introduce *Certified Selective Tracking* (CST), which delivers a result only when both deadline and branch-identity certificates hold. Its *Geometry-aware Certified Trackability* (GCT) instantiation encloses active-set events under an explicit uncertainty set. Strict separation of the event tubes certifies event membership, order, and endpoint identity. Within a finite candidate family, selection maximizes delivery availability under the same strict certificate contract; simultaneous envelope coverage bounds false-delivery risk. PTIF-KKT supplies inversion-free candidate updates. On 470 held-out PJM-driven arrivals in a constructed separable QP, public interval residual verification retains 134 deliveries at 5% curvature uncertainty with zero observed false identities. A true-model verifier retains 186, but 52 of these fail at an admissible curvature corner. Point-event selection makes 32 endpoint-label errors. A separate eight-dimensional heuristic measures agreement with offline multistart labels, not nonlinear branch certification or general solver superiority.

Keywords: time-varying optimization, KKT tracking, branch identification, interval verification, selective certification

1 INTRODUCTION

Constrained prediction heads, differentiable optimization layers, resource allocation modules, and safety filters repeatedly solve problems whose data and constraints change. The solver must reach a prescribed residual and feasibility tolerance before a deadline, even as the instance moves during computation. A small residual on a frozen instance alone cannot establish either timely arrival or the identity of the branch being tracked. We characterize when a moving constrained optimization problem can be tracked before a prescribed deadline through its Karush–Kuhn–Tucker (KKT) residual, while treating branch identity as a separate requirement. For a time-varying nonlinear program, the exact Fischer–Burmeister residual $E(y,t)$ combines stationarity, feasibility, and complementarity; y stacks primal and dual variables. Tracking before a deadline T_c requires a fixed-budget dynamics to bring $\|E(y(t),t)\|$ and the constraint violation below their targets before T_c.

Drift limits this ability: if $E(y,t)$ changes too quickly, local residual correction cannot keep pace within the deadline. Online optimization typically measures regret and cumulative violation, while static nonlinear programming studies convergence on a frozen instance. We instead require a *selective delivery contract*: CST reports a certified result only when it has evidence for both timely residual reduction and the declared branch identity, and otherwise abstains.

We use PTIF-KKT as the candidate-generation flow inside CST. Given a generalized Jacobian selection $G \in \partial_y E(y,t)$, PTIF-KKT evolves

$$\dot{y} = -\gamma G^\top \left(\Psi_\kappa(E, T_c) + \widehat{E}_t \right), \tag{1}$$

where Ψ_κ is a norm-scaled prescribed-time residual correction and $\widehat{E}_t$ estimates the explicit time derivative $E_t = \partial E/\partial t$. The update is inversion-free: it uses only $G^\top v$, not G^{-1}, a KKT factorization, or a Newton linear solve. The time-feedforward term compensates for the lag caused by changes in the optimality system during computation.

Reviewers: please read the Reviewer Guidelines (iclr.cc/Conferences/2027/ReviewerGuidelines) and the AI Policy for Reviewers (iclr.cc/Conferences/2027/AIPolicyForReviewers). If you used AI to expand, edit, or polish your review, please provide the input text to the LLM. Better still, consider skipping the LLM and submitting your original text: we, and the authors, are much more interested in your unedited thoughts than in what an LLM has to say. AI-assisted or not, you are putting your name and reputation behind your review: LLM-generated falsehoods, hallucinations or misrepresentations are subject to disciplinary action, which may include desk-rejecting all papers you have authored.

CST combines these established residual-dynamics components with an event-complete identity requirement: GCT must determine every admissible active-set event and its order before certifying delivery. The residual and identity conditions play distinct roles. For residual tracking, let $c(s)$ be the certified correction term and $d(s)$ the worst-case drift bound at residual magnitude s. Define the robust drift-domination margin

$$D_r(s) = c(s) - d(s). \tag{2}$$

When this margin is positive and its settling-time bound meets the deadline, PTIF-KKT is certified to reach the target residual in time. A separate directional condition identifies admissible drift that prevents residual decrease; failure of the conservative spectral margin alone does not imply this obstruction. The margin therefore identifies a trackability boundary within this residual-flow family.

Contributions. The paper makes three contributions.

1. We define CST/GCT as a two-bottleneck delivery contract that separates deadline trackability from target-branch identity and treats abstention as a valid output.
2. We prove drift-domination, event enclosure, robust identity ordering, finite-candidate maximality within a fixed certificate contract, and risk transfer, with an explicit diagonal/M-matrix versus generic-SPD boundary.
3. We evaluate the contract with PTIF-KKT candidate generation on held-out arrivals and examine its limits under nonconvex uncertainty. The separate high-dimensional heuristic measures empirical label agreement and supplies no nonlinear branch certificate.

2 Related Work

Time-varying optimization and residual dynamics address changing optimality systems; selective prediction formalizes abstention; interval verification supplies bounds valid over uncertainty sets (5; 6; 11; 20–23; 25; 26). CST combines these requirements in a delivery rule: both deadline trackability and event-complete identity must be certified.

The mathematical ingredients also connect to finite-time and fixed-time stability (3; 4), numerical solution of nonlinear programs (2; 24), and operator-splitting or differentiable optimization layers (16; 17; 20; 21). Interval arithmetic supplies enclosure semantics rather than a statistical confidence guarantee (25); selective and conformal prediction study abstention or marginal coverage under different assumptions (22; 23). CST keeps these roles separate by requiring both a deadline margin and an event-complete identity check before delivery.

Online and projection-free optimization provide additional baselines for time-varying constrained computation (7; 9); the solver implementations used for the numerical comparisons are documented by conic and interval-solver references (18; 27).

3 Deadline KKT Residual Tracking

3.1 Time-varying constrained problems

We consider a time-varying nonlinear program

$$\begin{aligned} \min_{x\in\mathbb{R}^n} \quad & f(x,t) \\ \text{s.t.} \quad & h(x,t) = 0, \qquad g(x,t) \le 0, \end{aligned} \tag{3}$$

where $h : \mathbb{R}^n \times \mathbb{R}_{\ge 0} \to \mathbb{R}^p$ and $g : \mathbb{R}^n \times \mathbb{R}_{\ge 0} \to \mathbb{R}^m$. The objective and constraints may be nonconvex in x; our target is residual tracking of KKT points, not global optimality. Let $y = (x, \lambda, \mu) \in \mathbb{R}^N$, $N = n + p + m$, stack primal variables and multipliers.

For inequality complementarity we use the Fischer–Burmeister map, denoted by ϕ_c:

$$\phi_c(a, b) = a + b - \sqrt{a^2 + b^2}, \tag{4}$$

which satisfies $\phi_c(a,b) = 0$ if and only if $a \geq 0$, $b \geq 0$, and $ab = 0$. With $a = \mu_i$ and $b = -g_i(x,t)$, define the exact KKT residual

$$E(y,t) = \begin{bmatrix} \nabla_x f(x,t) + \nabla_x h(x,t)\lambda + \nabla_x g(x,t)\mu \\ h(x,t) \\ \Phi_c(\mu, -g(x,t)) \end{bmatrix}. \tag{5}$$

At fixed t, $E(y,t) = 0$ encodes stationarity, equality feasibility, inequality feasibility, and complementarity under standard constraint qualifications. Because the FB map is nonsmooth at $(0,0)$, $E(\cdot,t)$ is generally locally Lipschitz rather than everywhere differentiable. We write $\partial_y E(y,t)$ for Clarke's generalized Jacobian.

Proposition 1 (Standard optimality meaning for affine QPs)**.** *Suppose that at a fixed t, (3) is a strongly convex quadratic program with affine constraints and that the FB-KKT residual has a local error bound*

$$\operatorname{dist}(y, \mathcal{Y}^\star(t)) \leq C_E \, \|E(y,t)\|$$

on the tube under consideration, where $\mathcal{Y}^\star(t)$ is the KKT solution set. If $\|E(y,t)\| \leq s$, then $\operatorname{dist}(y, \mathcal{Y}^\star(t)) \leq C_E s$, the primal distance to the optimizer is at most $C_E s$, and affine constraint violation is at most $L_c C_E s$, where L_c is the constraint-operator norm on the tube. For an affine-equality QP $\min_x \frac{1}{2} x^\top Q x + c^\top x$ subject to $Ax = b$, let $\Pi_A(x)$ be the Euclidean projection onto $Ax = b$. If $\Pi_A(x)$ stays in the tube, then

$$0 \leq f(\Pi_A(x), t) - f(x^\star, t) \leq \frac{\|Q\|}{2} C_E^2 s^2.$$

Thus a residual tube gives first-order distance and feasibility guarantees, and after equality repair it gives a second-order objective-gap certificate. In particular, $E(y,t) = 0$ gives a global optimizer of the convex QP.

Proof. The assumed local error bound gives the stated primal–dual distance estimate, and affine constraint violation is at most the constraint-operator norm times that distance. Euclidean projection onto $Ax = b$ is nonexpansive. Expanding the strongly convex quadratic around its equality-constrained minimizer cancels the feasible first-order term and leaves at most $\|Q\|\|\Pi_A(x) - x^\star\|^2/2$. KKT sufficiency then gives global optimality when $E(y,t) = 0$; see also (19; 24). □

This result translates the certified residual tolerance into distance, feasibility, and objective-gap bounds for the affine-QP experiments.

3.2 Deadline trackability

Definition 2 (Deadline residual target)**.** Given a deadline $T_c > 0$, a residual tolerance $\varepsilon_E \geq 0$, and a feasibility tolerance $\varepsilon_g \geq 0$, a solver achieves the deadline target from y_0 if

$$\|E(y(T_c), T_c)\| \leq \varepsilon_E, \qquad \max_i [g_i(x(T_c), T_c)]_+ \leq \varepsilon_g. \tag{6}$$

Definition 3 (Certified trackability)**.** An instance path of (3) is certified trackable on a residual range $\mathcal{S} = [s_\star, s_{\max}]$, $0 \leq s_\star < s_{\max}$, if a computable margin $D_r(s)$ is positive on $(s_\star, s_{\max}]$ and the residual trajectory remains in this range until it reaches $s_\star$. The case $s_\star = 0$ is an exact certificate; $s_\star > 0$ is a practical deadline certificate.

The residual range $\mathcal{S}$ can be chosen analytically, estimated from a warm start, or monitored online. Trackability is local to this range: a moving instance may be trackable from a warm start even when the certificate fails for more distant initializations.

4 PTIF-KKT: Inversion-Free Deadline Residual Dynamics

4.1 Norm-scaled prescribed-time correction

Fix $q \in (0,1)$, a scale $\kappa > 0$, and a deadline $T_c > 0$. For $E \neq 0$, set $s = \|E\|$ and define

$$\Psi_\kappa(E, T_c) = \kappa(T_c, q)\big(s^{-q} + s^q\big)E, \qquad \Psi_\kappa(0, T_c) = 0. \tag{7}$$

The correction is collinear with the residual. If $W \succeq \underline{\lambda} I$, then

$$E^\top W \Psi_\kappa(E, T_c) \geq \underline{\lambda}\, \kappa(T_c, q)\left(s^{2-q} + s^{2+q}\right). \tag{8}$$

4.2 CONTINUOUS DYNAMICS

PTIF-KKT uses the differential inclusion

$$\dot{y} \in -\gamma\, \partial_y E(y,t)^\top \left(\Psi_\kappa(E(y,t), T_c) + \widehat{E}_t(y,t) \right), \tag{9}$$

where $\gamma > 0$ and $\widehat{E}_t$ is an estimator of $E_t(y,t)$. Equivalently, along a measurable selection $G(t) \in \partial_y E(y(t),t)$,

$$\dot{y}(t) = -\gamma\, G(t)^\top \left(\Psi_\kappa(E(y(t),t), T_c) + \widehat{E}_t(y(t),t) \right). \tag{10}$$

The update is matrix-inverse-free. It requires residual evaluation, a generalized Jacobian selection or Jacobian-vector implementation, and the adjoint transport $G^\top v$.

4.3 CERTIFICATE-AWARE DIGITAL SOLVER

Algorithm 1: Evidence-Gated Event Delivery Router

Input: residual E, tracker state y_k, endpoint interface (arrived endpoint or declared forecast contract), estimator $\widehat{E}_t$, anchor $y_\star$, candidates $\mathcal{C}$, deadline T_c, range $\mathcal{S}$, error budgets, step h, optional compact domain X and feasible incumbent $\bar{x}$.

1. At tick $k{+}1$, record whether the complete endpoint has arrived. If it has, form E_{k+1}; apply event-causal affine transport when its algebraic contract holds, otherwise pass the tracker candidate to an endpoint re-anchor verifier.
2. If delivery is required before endpoint arrival, use a forecast update only under a declared, audited forecast envelope. Without that contract, emit Unresolved.
3. For a continuing tracker, evaluate each candidate's residual, derivative estimate, and D_r. If the endpoint is noisy or partial, add the declared finite-event radius $\varepsilon_o(\delta)$ to the estimator-error budget. Before committing a digital segment, enclose its full residual trace by U_k, compute $\underline{\delta}_b(r_k)$, and set $M_k = \underline{\delta}_b(r_k) - U_k$.
4. Convert the deadline and identity envelopes to $\vartheta_t = T_h^u / T_c$ and $\vartheta_i = \sup_t \bar{s}(t)/\delta_b(t, r(t))$. Set $\vartheta_i = +\infty$ unless the hypotheses of Theorem 9 are verified, including initial containment and isolation of the declared root throughout the tube.
   Set $\mathfrak{m}_* = \min\{1 - \vartheta_t, 1 - \vartheta_i\}$.
5. Deliver CertifiedIdentified (or CertifiedIdentifiedHP under the finite-event confidence contract) only for a candidate with $\mathfrak{m}_* > 0$; choose the largest two-bottleneck slack.
6. If branch identity fails, interval-prune separated inequalities and classify every remaining local chart as Krawczyk-enclosed, excluded, degenerate, or unresolved. If an arrived endpoint has a unique valid local chart, emit VerifiedReanchor.
7. If anchor coverage is unknown, search the compact domain over all LICQ-sized charts; each box must be objective-pruned, root-excluded, uniquely enclosed, or unresolved. Commit only after the complete tree closes and one valid root uniformly dominates.
8. Otherwise emit Unresolved; log interface type, residual, both ratios, two-bottleneck slack, recovery certificate, and cost.

Both identified statuses require the hypotheses of Theorem 9, including initial containment and an audited isolating tube. The high-probability status additionally requires Theorem 14; confidence coverage does not replace root isolation. Subject to these prerequisites, each output status specifies which properties have been established:

Status	Required evidence	Reported meaning
CertifiedIdentified	residual margin and branch margin positive	same-branch result delivered
ResidualOnly	residual target met but branch margin fails	numerical result; identity withheld
Unresolved	failed residual or missing audit	no semantic result delivered
VerifiedReanchor	arrived endpoint; local-root tests pass	endpoint-local anchor
CertifiedGlobal	complete target tree and strict winner	unique global target delivered

When a margin is unknown or nonpositive, tracking may continue, but branch identity remains unresolved and recovery is required before certified delivery. Event-causal and anticipatory delivery use different information sets; their latency and tracking claims are evaluated separately.

4.4 GEOMETRY-AWARE CERTIFIED TRACKABILITY

Geometry-aware Certified Trackability (GCT) implements the event-complete identity requirement for a piecewise-affine KKT path. It follows the public path, encloses all possible active-set events in

a declared curvature box, and accounts for the cost of each event update. For a separable box QP, let $T_i^L = [L_i^-, L_i^+]$ and $T_i^U = [U_i^-, U_i^+]$ denote the interval tubes for lower/free and free/upper events. GCT computes an observable geometry margin

$$\mu_g = \min\{\mu_e, \mu_p, \mu_o\},$$

where the three terms are respectively the distance of the endpoint neighborhood from every event tube, the distance of non-crossed tubes from the transition path, and the gap between consecutive crossed tubes. A zero term means that the corresponding identity fact is not determined.

Given a correction budget B, event count K, and independent verifier cost V, GCT returns CertifiedTrackable only if

$$\mu_g > 0, \qquad \|E\| \le \varepsilon_E, \qquad 1 + K + V \le B.$$

If endpoint identity is robust but path ordering is unresolved, it returns EndpointOnly; if only the residual gate passes, it returns ResidualOnly; otherwise it abstains. The target active-set label is unavailable to the selector. Event prediction, geometry, and numerical accuracy are evaluated separately, so a small residual cannot substitute for missing identity evidence.

Theorem 4 (GCT branch-budget soundness)**.** *Assume positive diagonal curvature, a valid declared curvature box, exact offsets and capacities, strict event-tube separation, and a terminal residual verifier valid uniformly over that box for the emitted candidate. If GCT returns* CertifiedTrackable*, then every admissible curvature in the box has the same crossed event set and event order, the endpoint has the declared active-set identity, the emitted candidate satisfies the residual target, and the correction plus verification cost is at most B.*

Proof. The interval event enclosure contains every admissible switch. Strict path separation fixes event membership and pairwise disjointness fixes event order. The KKT solution is affine between consecutive switches, so induction over the ordered event list preserves event identity and order. The endpoint tube gate excludes competing active sets, while the residual and budget gates give the deadline contract. □

Overlapping tubes require abstention. Public interval residual verification retains 445/470 held-out deliveries at zero curvature mismatch and 134/470 at 5%, with zero observed false identities. A verifier using the realized curvature gives the less restrictive 186/470 operating point.

5 Trackability Theory

5.1 Residual derivative

Assumption 5 (Regular residual arc)**.** On an open tube containing the analyzed trajectory, E is locally Lipschitz jointly in (y, t), and its explicit derivative E_t exists a.e. along the trajectory. The measurable selection $G(t) \in \partial_y E(y(t), t)$ used by the dynamics must satisfy the chain rule

$$\dot{E}(y(t), t) = G(t)\dot{y}(t) + E_t(y(t), t) \tag{11}$$

holds for almost every t.

Lemma 6 (Residual evolution)**.** *Under Assumption 5, the selection dynamics (10) give, almost everywhere,*

$$\dot{E} = E_t - \gamma W \left(\Psi_\kappa(E, T_c) + \widehat{E}_t\right), \qquad W := GG^\top. \tag{12}$$

Consequently, for $V = \frac{1}{2}\|E\|^2$,

$$\dot{V} = E^\top E_t - \gamma E^\top W \Psi_\kappa - \gamma E^\top W \widehat{E}_t. \tag{13}$$

Proof. Substitute (10) into (11); the identity for V follows from absolute continuity. □

5.2 Robust drift-domination margin

Let $s = \|E\|$. The certificate compares the coercive residual correction with all terms that can increase V.

Definition 7 (Certified residual envelope)**.** A certified residual envelope on a tube $\mathcal{S} = [s_\star, s_{\max}]$ consists of bounds $(\underline{\lambda}, \bar{\lambda}, d, \varepsilon_t, R_e)$ such that, for the selected generalized-Jacobian path and implemented vector field, whenever $s(t) = \|E(y(t), t)\| \in \mathcal{S}$,

$$\underline{\lambda} I \preceq W(t) \preceq \bar{\lambda} I, \qquad \left\| \widehat{E}_t - E_t \right\| \le \varepsilon_t, \tag{14}$$

$$\|E_t(y(t), t)\| \le d(s(t)), \qquad d(s) := \eta_1 + \eta_2 s \tag{15}$$

and implementation effects such as stale Jacobians, clipping, and discretization are upper bounded by a nonnegative Lyapunov-defect budget $R_e(s)$. Analytic or interval-audited bounds may support a Certified label; empirical estimates are diagnostics only.

The bounds in Definition 7 must be certified on $\mathcal{S}$; otherwise Algorithm 1 reports Uncertified. Define the feedforward mismatch factor

$$\chi_\gamma := \sup_t \|I - \gamma W(t)\| \le \max_{\lambda \in [\underline{\lambda}, \bar{\lambda}]} |1 - \gamma \lambda|. \tag{16}$$

Define

$$\begin{aligned} D_r(s) := \gamma \underline{\lambda} \kappa(T_c, q) \left(s^{2-q} + s^{2+q} \right) \\ - \chi_\gamma s d(s) - \gamma \bar{\lambda} \varepsilon_t s - R_e(s). \end{aligned} \tag{17}$$

The form of R_e depends on the digital implementation; Corollary 12 gives a common specialization. Operationally, R_e is a one-step residual-Lyapunov defect budget: a discrete run may inherit the continuous certificate only when its logged or bounded defect is subtracted from the margin before the status is assigned. For a target radius $s_\star \ge 0$, the normalized margin over $\mathcal{S} \cap (s_\star, \infty)$ is

$$\rho_{\mathcal{S}} := \inf_{s \in \mathcal{S}:\, s > s_\star} \frac{D_r(s)}{s^{2-q} + s^{2+q}}. \tag{18}$$

The case $\rho_{\mathcal{S}} > 0$ is the certified trackable regime down to $s_\star$. If residual-independent drift or estimator error remains near zero, usually only $s_\star > 0$ can satisfy this condition.

Constructing the envelopes. Certification requires bounds over the tube and time window traversed by the solver. For the selected G,

$$\underline{\lambda} = \inf_{\mathcal{S}} \sigma_{\min}^2(G), \qquad \bar{\lambda} = \sup_{\mathcal{S}} \sigma_{\max}^2(G),$$

over the monitored tube and time window; drift, derivative-error, stale-Jacobian, and integration budgets come from analytic, automatic-differentiation, interval, or inflated sampled audits. Analytic or interval bounds support a Certified output; sampled estimates provide empirical diagnostics even when the numerical target is reached.

Theorem 8 (Drift-domination certificate)**.** *Let $\mathcal{S} = [s_\star, s_{\max}]$, $0 \le s_\star < s_{\max}$. Let Assumption 5 and (14)–(15) hold along a PTIF-KKT trajectory while $s(t) \in (s_\star, s_{\max}]$. Suppose $\rho_{\mathcal{S}} > 0$. For any split $\theta \in (0, 1)$, define*

$$\alpha_\theta := \theta \rho_{\mathcal{S}} \, 2^{\frac{2-q}{2}}, \qquad \beta_\theta := (1 - \theta) \rho_{\mathcal{S}} \, 2^{\frac{2+q}{2}}. \tag{19}$$

Then, for almost every t with $s(t) \in (s_\star, s_{\max}]$,

$$\dot{V} \le -\alpha_\theta V^{\frac{2-q}{2}} - \beta_\theta V^{\frac{2+q}{2}}. \tag{20}$$

Let $V_\star = \frac{1}{2} s_\star^2$. If the trajectory remains in $(s_\star, s_{\max}]$ until reaching the target tube, then V reaches $V_\star$ no later than

$$\begin{aligned} T_h &\le \frac{2}{q \sqrt{\alpha_\theta \beta_\theta}} \left[\arctan\left(V(0)^{q/2} \sqrt{\frac{\beta_\theta}{\alpha_\theta}} \right) - \arctan\left(V_\star^{q/2} \sqrt{\frac{\beta_\theta}{\alpha_\theta}} \right) \right] \\ &\le \frac{2}{q \sqrt{\alpha_\theta \beta_\theta}} \arctan\left(V(0)^{q/2} \sqrt{\frac{\beta_\theta}{\alpha_\theta}} \right) \\ &\le \frac{\pi}{q \sqrt{\alpha_\theta \beta_\theta}}. \end{aligned} \tag{21}$$

If $\pi/(q\sqrt{\alpha_\theta \beta_\theta}) \le T_c$, the residual is certified to enter the radius-$s_\star$ tube by the deadline. Exact convergence corresponds to the special case $s_\star = 0$ with a positive margin down to zero.

The complete proof is in Appendix A.3.

Appendix C.5 derives the spectral initializer and the conditions under which gain selection preserves this certificate.

5.3 Minimax boundary

The local minimax result gives the exact directional condition under which admissible drift can prevent residual decrease; its statement and proof are in Appendix A.4.

5.4 Branch identity in nonconvex KKT paths

For nonconvex problems, a small KKT residual does not by itself identify which stationary branch is being tracked. We therefore separate the residual certificate from a local branch-identity certificate. For a declared branch $y_\star(t)$ and a continuous tube radius $r(t)$, define the residual escape barrier $\delta_b(t, r(t)) = \inf_{\|y - y_\star(t)\| = r(t)} \|E(y,t)\|$ and the margin $M_i(t) = \delta_b(t, r(t)) - \bar{s}(t)$, where $\bar{s}(t)$ is the audited residual envelope. The moving-tube version of Theorem 18 proves that positive margin prevents a first exit. For an isolating tube, D_r certifies residual tracking, while M_i certifies that the trajectory remains in the declared nonconvex KKT branch tube.

The router uses both margins to select its output. It delivers a same-branch result only when both are positive; otherwise it reports residual-only or re-anchor status.

The nonconvex diagnostic uses the realizable quartic $f(x,t) = x^4/4 - ax^2/2 - b(t)x$, whose minimum and saddle branches merge at a fold; Appendix C gives the formal certificate, derivation, and phase diagram.

For nonlinear inequalities, interval slack pruning reduces 2^m active sets to 2^{k_r} leaves on anchor tube r, where k_r is its local ambiguity width. Theorem 28 proves sound identity relative to covered tubes. If that coverage is unknown, Theorem 30 starts from a compact domain and verified incumbent, enumerates the $\sum_{j=0}^{n} \binom{m}{j}$ LICQ-sized charts, and permits delivery only after the target-relevant interval tree closes. Thus coverage is proved or explicitly withheld, but worst-case spatial search and active-chart growth remain exponential (Appendices C.3 and C.4).

Digital certification. Clipping, stale Jacobians, integration defect, and affine repair preserve a Certified label only when their a priori envelopes are included in the digital drift margin; logged defects alone remain empirical. The exact digital corollary, observed-defect definition, and repair proposition are in Appendix A.1.

5.5 The two-bottleneck trackability certificate

To combine the drift and identity conditions, first express each as a dimensionless ratio. Define the deadline ratio

$$\vartheta_t := \frac{T_h^u}{T_c}, \tag{22}$$

where T_h^u is the upper bound in Theorem 8. For a declared branch and an audited residual envelope $\bar{s}(t)$, define the identity ratio

$$\vartheta_i := \sup_{t \in [0, T_c]} \frac{\bar{s}(t)}{\delta_b(t, r(t))}, \tag{23}$$

with value $+\infty$ when the denominator vanishes. The two-bottleneck certificate is

$$\mathfrak{m}_* := \min\{1 - \vartheta_t,\, 1 - \vartheta_i\}. \tag{24}$$

Taking the minimum requires both normalized margins to be positive without adding quantities with different physical units.

Theorem 9 (Two-bottleneck identified trackability)**.** *Assume the regular residual arc, a certified residual envelope, and a continuous declared KKT branch $y_\star(t)$ with an audited escape barrier on*

a moving tube $r(t)$ containing no other KKT root at any time. Suppose the initial state lies strictly inside that tube, the drift-domination hypotheses of Theorem 8 hold, and the residual envelope used in (23) bounds the entire trajectory until the deadline. If $\mathfrak{m}_* > 0$, then the PTIF-KKT trajectory reaches the declared residual target by T_c and remains in the declared KKT branch tube throughout $[0, T_c]$. Thus the output has both deadline residual and same-branch identity semantics.

If $\vartheta_t \geq 1$, the deadline claim is withheld. If $\vartheta_i \geq 1$ or the barrier is unaudited, the residual target may still be reported as ResidualOnly*, but branch identity is withheld. This is an unresolved identity condition, not a claim that the trajectory is on the wrong branch.*

Proof. Positive $1 - \vartheta_t$ gives $T_h^u < T_c$, so Theorem 8 reaches the residual target. Positive $1 - \vartheta_i$ is equivalent to $\bar{s}(t) < \delta_b(t, r(t))$ for every t. The first-exit argument in Theorem 18 therefore prevents the trajectory from leaving the declared tube. If either strict inequality fails, the corresponding implication is unavailable, which is exactly the selective output rule. □

Certified delivery requires the deadline and identity conditions to hold together. The phase-diagram experiments report each condition separately and their intersection, exposing which requirement limits delivery. Figure 2 in the appendix depicts this strict intersection and the event-order ambiguity that the identity gate must resolve.

Proposition 10 (Maximal delivery within a fixed strict certificate contract)**.** *Fix a finite candidate family $\mathcal{C}$ and its audited evidence. Let $\mathcal{D} = \{c : \vartheta_t(c) < 1\}$ and $\mathcal{I} = \{c : \vartheta_i(c) < 1\}$, with missing audits assigned ratio $+\infty$. Restrict admissible selectors to delivering only from $\mathcal{D} \cap \mathcal{I}$. The maximum-slack two-bottleneck selector delivers if and only if this intersection is nonempty, and thus maximizes delivery availability within this contract. Under the soundness hypotheses of Theorem 9, its delivered result is sound. This is not a coverage bound for selectors using other evidence or weaker sufficient conditions. In particular, deadline equality is excluded by this strict contract even when Theorem 8 certifies arrival at T_c.*

Proof. If the intersection is empty, the contract forbids every delivery. Otherwise, finiteness ensures a maximum-slack member exists with $\mathfrak{m}_* > 0$; selecting it delivers whenever any contract-admissible selector could. Soundness follows from Theorem 9, not from a necessity claim about these bounds. □

Uncertainty and false-delivery risk. CST accepts only a declared envelope. Let δ_e denote its coverage-failure probability and δ_n the numerical-verifier failure probability. If the envelope has simultaneous coverage $1 - \delta_e$ over the registered horizon and the numerical verifier fails with probability at most δ_n, false identified delivery has probability at most $\delta_e + \delta_n$, without an event- independence assumption. Analytic/interval bounds set $\delta_e = 0$; statistical envelopes need simultaneous rather than pointwise coverage. Without either contract CST abstains. The formal transfer and a stochastic construction are in Appendix A.2. The primary experiment below is conditional on its declared mismatch boxes.

Definition 11 (Budgeted CST evidence graph)**.** For a finite candidate family, an evidence graph has auditable state z, remaining integer budget b, actions with costs, and terminal reward one only when both deadline and identity gates pass. Hidden endpoint and branch labels are not graph observations.

Appendix A.2 states the corresponding finite-event high-probability result. Its scope is the digital observation interface; general nonlinear convergence requires separate analysis.

6 Experiments

We evaluate CST on 470 chronological held-out PJM load arrivals driving a constructed six-variable separable QP. Curvature is uncertain; offsets and capacities are exact. The public residual gate bounds the emitted candidate over the curvature box, without reading the realized curvature. Terminal identity is reserved for evaluation. Figure 1 separates the requirements for delivery; Table 1 in the appendix retains the exact counts. Point-event selection always reports its predicted identity. The endpoint-tube rule requires robust endpoint identity, the complete-tube rule also requires event

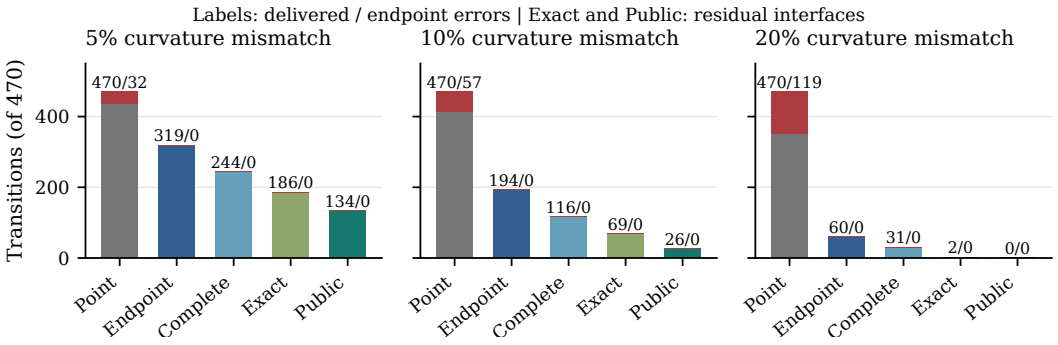


Figure 1: Held-out delivery under declared curvature mismatch. Each bar splits delivered transitions into correct endpoint labels (solid) and errors (hatched); labels give delivered/error counts. Endpoint and complete tubes impose successively stronger identity requirements. Exact and public CST use realized-model and uniform-box residual checks with the same event budget. Zero observed errors do not establish unconditional zero risk. All panels use the same 470 held-out arrivals.

membership and order, and CST additionally checks residual and budget conditions. Each additional requirement can reduce coverage while withholding claims that the available evidence cannot certify.

At 5% mismatch, the true-model residual gate accepts 186 arrivals, but 52 of their candidates violate the residual target at a permitted curvature corner. The public interval gate retains 134; at 10% and 20% it retains 26 and zero. These are conditional QP replay results, not hardware evidence. Appendix C.6 gives the public bound and independent constructed-demand controls, and distinguishes the dense-Hessian and nonconvex extensions.

Exact arithmetic check for the public residual gate. The public residual audit uses a saved-binary64 contract: each archived candidate, demand, multiplier, curvature endpoint, offset, and capacity is treated as the exact dyadic rational represented by its binary64 value. Using Python `Fraction` arithmetic, we recomputed the squared residual, including clipping and the coordinate-wise maximum, for all 1,880 rows (470 arrivals at each of four curvature levels). The independent 64-corner enumeration agreed on all 120,320 corners. With the exact decimal threshold $0.02 = 1/50$ and the residual predicate $R^2 \le (1/50)^2$, the delivered counts were 445, 134, 26, and 0 at 0, 5, 10, and 20% uncertainty; they matched the saved Float64 decisions in every row. The minimum delivered squared-residual slack at 5% was positive (approximately 1.390×10^{-6}); the closest all-row slack was negative, so the audit does not turn rejected rows into deliveries. The binary64 literal `0.02` is slightly larger than $1/50$; replacing the decimal threshold by that literal changed no row. This result is limited to arithmetic for the fixed saved candidates and does not enclose upstream source-real inputs or certify event geometry.

7 Conclusion

Certified Selective Tracking delivers a result when deadline and branch-identity evidence hold together. On the held-out transitions, this rule yields zero observed false identities under declared mismatch, with decreasing coverage as uncertainty grows. Delivery availability is maximal only within the same strict certificate contract. The bounded nonconvex study and the observed-only high-dimensional heuristic identify distinct limits of the evidence; the latter supplies no branch certificate.

Reproducibility Statement

The main text and appendix state the assumptions and proof details for the certificate claims. The supplementary material contains the bounded extension tables, while the accompanying artifact package provides the scripts, frozen results, manifests, and verification command needed to repro-

duce the reported results. The high-dimensional protocol and results have a separate version from the historical CSVs; reference labels are used only for scoring.

AI USE STATEMENT

Generative AI tools assisted with mathematical revisions, experimental design, implementation, and analysis of results. The authors take responsibility for the final content.

REFERENCES

[1] F. H. Clarke. *Optimization and Nonsmooth Analysis*. SIAM, 1990.

[2] A. Fischer. A special Newton-type optimization method. *Optimization*, 24(3–4):269–284, 1992.

[3] S. P. Bhat and D. S. Bernstein. Finite-time stability of continuous autonomous systems. *SIAM Journal on Control and Optimization*, 38(3):751–766, 2000.

[4] A. Polyakov. Nonlinear feedback design for fixed-time stabilization of linear control systems. *IEEE Transactions on Automatic Control*, 57(8):2106–2110, 2012.

[5] E. Dall'Anese and A. Simonetto. Prediction-correction algorithms for time-varying constrained optimization. *IEEE Transactions on Signal Processing*, 65(20):5481–5494, 2017.

[6] A. Simonetto, E. Dall'Anese, S. Paternain, G. Leus, and G. B. Giannakis. Time-varying convex optimization: Time-structured algorithms and applications. *Proceedings of the IEEE*, 108(11):2032–2048, 2020.

[7] M. Rostami and S. S. Kia. Time-varying convex optimization with $O(n)$ computational complexity. *Automatica*, 178:112355, 2025.

[8] T. Kamijima, N. Marumo, and A. Takeda. A simple yet highly accurate prediction-correction algorithm for time-varying optimization. *arXiv preprint arXiv:2504.05798*, 2025.

[9] D. Garber and B. Kretzu. Projection-free online convex optimization with time-varying constraints. In *International Conference on Machine Learning*, 2024.

[10] S. Hutchinson, T. Chen, and M. Alizadeh. Optimistic safety for online convex optimization with unknown linear constraints. In *International Conference on Artificial Intelligence and Statistics*, 2025.

[11] J. Chen, Y. Pan, Y. Zhang, S. Li, and N. Tan. Inverse-free zeroing neural network for time-variant nonlinear optimization with manipulator applications. *Neural Networks*, 178:106462, 2024.

[12] Z. Zhang, X. Sun, X. Li, and Y. Liu. An adaptive variable-parameter dynamic learning network for solving constrained time-varying QP problem. *Neural Networks*, 184:106968, 2025.

[13] Y. Zeng, C. Hua, B. Liao, and Z. Li. Design and analysis of gradient-based differential neural network for solving time-varying quadratic problems with inequality constraint. *Neural Networks*, 192:107911, 2025.

[14] G. Cheng, B. Qiu, J. Guo, and Y. Han. A robust direct-discretized RNN for time-dependent optimization constrained by nonlinear equalities and its applications. *IEEE/CAA Journal of Automatica Sinica*, 12(9):1866–1877, 2025.

[15] Z. Zhang, X. Sun, Y. Liu, and Y. Luo. A meta-interactive neural network for solving time-varying quadratic programming problems. *Nature Communications*, 16:10284, 2025.

[16] B. Stellato, G. Banjac, P. Goulart, A. Bemporad, and S. Boyd. OSQP: An operator splitting solver for quadratic programs. *Mathematical Programming Computation*, 12:637–672, 2020.

[17] B. O'Donoghue, E. Chu, N. Parikh, and S. Boyd. Conic optimization via operator splitting and homogeneous self-dual embedding. *Journal of Optimization Theory and Applications*, 169(3):1042–1068, 2016.

[18] P. Goulart and Y. Chen. Clarabel: An interior-point solver for conic programs with quadratic objectives. *arXiv preprint arXiv:2405.12762*, 2024.

[19] S. Boyd and L. Vandenberghe. *Convex Optimization*. Cambridge University Press, 2004.

[20] B. Amos and J. Z. Kolter. OptNet: Differentiable optimization as a layer in neural networks. In *International Conference on Machine Learning*, 2017.

[21] A. Agrawal, B. Amos, S. Barratt, S. Boyd, S. Diamond, and J. Z. Kolter. Differentiable convex optimization layers. In *Advances in Neural Information Processing Systems*, 2019.

[22] Y. Geifman and R. El-Yaniv. Selective classification for deep neural networks. In *Advances in Neural Information Processing Systems*, 2017.

[23] A. N. Angelopoulos and S. Bates. Conformal prediction: A gentle introduction. *Foundations and Trends in Machine Learning*, 16(4):494–591, 2023.

[24] J. Nocedal and S. J. Wright. *Numerical Optimization*. Springer, 2006.

[25] R. E. Moore, R. B. Kearfott, and M. J. Cloud. *Introduction to Interval Analysis*. SIAM, 2009.

[26] I. Araya, F. Messine, J. Ninin, and G. Trombettoni. Hybridizing two linear relaxation techniques in interval-based solvers. *Journal of Global Optimization*, 91:437–456, 2025.

[27] H. Zhang, T. Kerkenhoff, N. Kichler, M. Dahmen, A. Mitsos, U. Naumann, and D. Bongartz. Accelerating deterministic global optimization via GPU-parallel interval arithmetic. *arXiv preprint arXiv:2507.20769*, 2025.

[28] R. Prakash, S. Janardhanan, and S. Sen. Computational complexity analysis of interval methods in solving uncertain nonlinear systems. *arXiv preprint arXiv:2603.19965*, 2026.

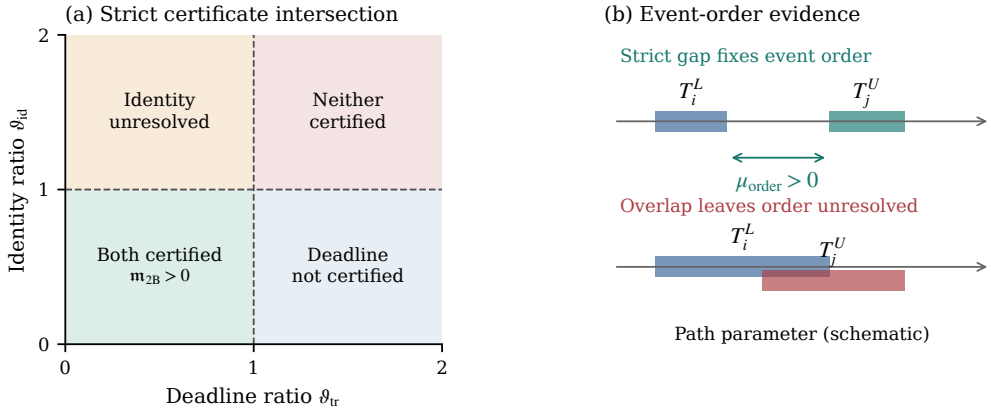


Figure 2: CST's two-bottleneck contract and event-order evidence. The certified region requires both normalized ratios to be below one, so the two-bottleneck margin is positive. Separated event tubes determine order; overlapping tubes leave identity unresolved. The event-tube panel is a schematic rendering of the interval definitions, not an experimental trace.

Table 1: Exact held-out counts underlying Figure 1. Each entry gives delivered/endpoint-label errors out of 470. Exact and public CST use realized-model and uniform-box residual verification, respectively.

mismatch	point event	endpoint	complete	CST exact	CST public
0%	470/0	445/0	445/0	445/0	445/0
5%	470/32	319/0	244/0	186/0	134/0
10%	470/57	194/0	116/0	69/0	26/0
20%	470/119	60/0	31/0	2/0	0/0

Appendix Guide

The appendix develops the digital and probabilistic certificates, the nonsmooth residual analysis, and the nonconvex branch conditions. It concludes with the structural limits of event-tube certification. Implementation logs, exploratory diagnostics, and extended workload tables accompany the code.

A Proof Details

A.1 Digital implementation certificates

Corollary 12 (Practical digital certificate)**.** *Consider an absolutely continuous interpolant of a digital implementation whose radial clipping satisfies $\dot{y} = \beta(t)\dot{y}_0$, with $0 < \beta_{\min} \le \beta(t) \le 1$. Assume stale-Jacobian mismatch is bounded by ζ_G, and one-step integration slack contributes at most $c_h h^r s$ to $\dot{V}$ almost everywhere on each segment in $\mathcal{S}$. In particular, the implemented step must satisfy*

$$\frac{V_{k+1} - V_k}{h} \le \frac{1}{h} \int_{t_k}^{t_{k+1}} \left(\dot{V}_c(t) + c_h h^r s(t) \right) dt$$

whenever the status is reported as certified. An endpoint defect inequality alone is not equivalent to this interpolation assumption and supplies no intersample guarantee. Define

$$\chi_{\gamma,\beta} := \sup_t \|I - \gamma\beta(t)W(t)\| \le \max_{\lambda \in [\underline{\lambda}, \bar{\lambda}],\, \beta \in [\beta_{\min}, 1]} |1 - \gamma\beta\lambda|. \tag{25}$$

Then Theorem 8 applies after replacing the correction and drift terms in (17) by

$$D_r^d(s) := \gamma\beta_{\min}\underline{\lambda}\kappa(T_c, q)\left(s^{2-q} + s^{2+q}\right) - \chi_{\gamma,\beta} s d(s) - \gamma\bar{\lambda}\varepsilon_t s - R_e(s), \tag{26}$$

where

$$R_e(s) = c_\zeta \zeta_G \, s \left[\kappa(T_c, q)(s^{1-q} + s^{1+q}) + d(s) + \varepsilon_t\right] + c_h h^r s \tag{27}$$

is one common implementation slack. If the defect inequality is unavailable, the same iterate can still be reported, but the status is downgraded to Practical *or* Uncertified.

Proof. Apply Theorem 8 to the clipped residual dynamics with $\chi_{\gamma,\beta}$, $\beta_{\min}$, and the stated implementation slack. □

A.2 High-probability interface and polynomial-budget limits

$$\varepsilon_o(\delta) := b + \sigma\left(\sqrt{d_o} + \sqrt{2\log(K/\delta)}\right), \tag{28}$$

Proposition 13 (Envelope-to-delivery risk transfer)**.** *For a predeclared horizon of K events, let $\widehat{\Theta}$ be an uncertainty envelope constructed without terminal identity labels and let $\theta_{1:K}$ denote the realized parameters. Suppose CST is sound whenever $\theta_{1:K} \in \widehat{\Theta}$ and $\Pr(\theta_{1:K} \in \widehat{\Theta}) \geq 1 - \delta_e$. Then the probability of any false identified delivery over the horizon is at most δ_e. With simultaneous numerical-verifier failure probability δ_n, it is at most $\delta_e + \delta_n$. No independence between events is needed.*

Proof. On simultaneous containment and verifier success, CST soundness excludes a false identified delivery. The false-delivery event is therefore contained in the union of the two contract failures. □

Theorem 14 (Finite-event high-probability two-bottleneck certificate)**.** *Consider K predeclared digital events. At event k, let*

$$\widehat{E}_{t,k} = P_k E_{t,k} + \xi_k, \qquad \|(I - P_k)E_{t,k}\| \leq b,$$

where P_k is a past-measurable orthogonal projector of rank $r_k \leq d_o$. Conditional on the past, require $\xi_k = U_k z_k$, where the past-measurable columns of U_k form an orthonormal basis of $\mathrm{range}(P_k)$ and $z_k \sim \mathcal{N}(0, \sigma^2 I_{r_k})$. Thus noise has no component outside this at-most-d_o-dimensional subspace; rank zero implies zero noise. No independence across events is required. Let $\varepsilon_o(\delta)$ be defined by (28). Replacing ε_t by $\varepsilon_t + \varepsilon_o(\delta)$ in D_r^d and propagating the resulting envelope through all K segments gives ratios ϑ_t^p and ϑ_i^p. If both are below one, then with probability at least $1 - \delta$, every delivered segment reaches its declared residual target and stays in the declared branch tube. This is a finite digital event-interface statement, not a continuous sample-path claim.

Proof. Since $\|\xi_k\| = \|z_k\|$ and $r_k \leq d_o$, conditional Gaussian norm concentration gives a failure probability at most δ/K for each event at the radius in (28); a union bound gives simultaneous validity with probability $1 - \delta$. On that event the enlarged estimator-error term is an a priori bound in D_r^d. Applying Corollary 12 segmentwise and then Theorem 9 proves the claim. □

$$\chi_B(\mu, L) = \frac{1}{\left|T_B\left(-\frac{L+\mu}{L-\mu}\right)\right|}. \tag{29}$$

Proposition 15 (Residual-polynomial budget lower bound)**.** *For the inversion-free class with a polynomial fixed independently of the realized spectrum and initial residual,*

$$y_B = y_0 - G^\top p_{B-1}(GG^\top)E_0, \qquad p_{B-1} \in \Pi_{B-1},$$

and the spectral family $\mu I \preceq GG^\top \preceq LI$, no budget-$B$ member can uniformly contract every initial residual below $\chi_B(\mu, L)$ in (29). Equality is attained by the scaled Chebyshev residual polynomial. The statement is restricted to this fixed-polynomial interface and does not lower-bound spectrum-adaptive Krylov methods, arbitrary nonlinear solvers, or a fixed finite spectrum.

Proof. Writing $r_B(\lambda) = 1 - \lambda p_{B-1}(\lambda)$ gives $r_B \in \Pi_B$ and $r_B(0) = 1$. The Chebyshev alternation theorem solves $\inf_{r_B} \max_{[\mu,L]} |r_B|$, yielding $1/|T_B(-(L+\mu)/(L-\mu))| = \chi_B(\mu, L)$. □

A.3 Drift-domination certificate

From Lemma 6, add and subtract $\widehat{E}_t - E_t$ and use Cauchy–Schwarz:

$$\dot{V} \le -\gamma E^\top W \Psi_\kappa + \|E\| \, \|I - \gamma W\| \, \|E_t\| + \gamma \, \|E\| \, \|W\| \left\| \widehat{E}_t - E_t \right\| + R_e(s).$$

Using (14), (15), and (8) yields $\dot{V} \le -D_r(s)$. By (18), $D_r(s) \ge \rho_{\mathcal{S}}(s^{2-q} + s^{2+q})$. Splitting the two powers by θ and using $s^{2\pm q} = (2V)^{(2\pm q)/2}$ gives (20).

Let $p = (2-q)/2 \in (1/2, 1)$, $r = (2+q)/2 \in (1, 3/2)$, and suppose $\dot{V} \le -\alpha V^p - \beta V^r$. Then

$$T_h \le \int_{V_\star}^{V(0)} \frac{dz}{\alpha z^p + \beta z^r} = \frac{2}{q\sqrt{\alpha\beta}} \left[\arctan\left(z^{q/2} \sqrt{\frac{\beta}{\alpha}} \right) \right]_{V_\star}^{V(0)} \le \frac{\pi}{q\sqrt{\alpha\beta}}.$$

A.4 Minimax boundary

Proposition 16 (Residual-level minimax boundary)**.** *Fix a residual radius $s_0 > 0$, a residual vector E_0 with $\|E_0\| = s_0$, and a symmetric positive semidefinite matrix W. Consider the exact-feedforward residual dynamics generated by a collinear correction*

$$\dot{E} = -\gamma a(s) W E + (I - \gamma W) E_t, \qquad a(s) \ge 0, \tag{30}$$

which is (12) with $\widehat{E}_t = E_t$ and $\Psi = a(s)E$. For the admissible residual-drift ball $\mathcal{D}(s_0) = \{v : \|v\| \le d(s_0)\}$,

$$\sup_{E_t \in \mathcal{D}(s_0)} \dot{V} = -\gamma a(s_0) E_0^\top W E_0 + d(s_0) \, \|(I - \gamma W) E_0\| .$$

Consequently, if

$$\gamma a(s_0) E_0^\top W E_0 \le d(s_0) \, \|(I - \gamma W) E_0\| , \tag{31}$$

then an admissible residual drift in the declared envelope makes $\dot{V} \ge 0$ at $E = E_0$. Thus the directional supremum is exact at the residual-envelope level; the spectral lower margin in (17) can remain conservative.

Proof. For $V = \frac{1}{2}\|E\|^2$, (30) gives

$$\dot{V} = -\gamma a(s_0) E_0^\top W E_0 + E_0^\top (I - \gamma W) E_t.$$

The supremum of the second term over $\|E_t\| \le d(s_0)$ is

$$d(s_0) \, \left\| (I - \gamma W)^\top E_0 \right\| = d(s_0) \, \|(I - \gamma W) E_0\| ,$$

because W is symmetric. The maximizing drift is any vector collinear with $(I - \gamma W)E_0$, with norm $d(s_0)$. This proves the displayed supremum; if (31) holds, the same maximizing drift gives $\dot{V} \ge 0$. □

A.5 Digital certificate

With clipping, the residual derivative becomes

$$\dot{E} = E_t - \gamma \beta(t) W (\Psi_\kappa + \widehat{E}_t) + \delta_e,$$

where δ_e denotes stale-Jacobian and integration errors. The correction term is bounded below by

$$\gamma \beta_{\min} \underline{\lambda} \kappa(T_c, q)(s^{2-q} + s^{2+q}),$$

while the drift channel is bounded by

$$s \, \|I - \gamma \beta(t) W\| \, \|E_t\| \le \chi_{\gamma,\beta} s d(s).$$

The estimator error contributes at most $\gamma \bar{\lambda} \varepsilon_t s$, since $\beta(t) \le 1$. The stated R_e upper bounds the remaining implementation terms in $\dot{V}$, so Theorem 8 applies with D_r^d in place of D_r.

B Nonsmooth Residual Details

The scalar Fischer–Burmeister map

$$\phi_c(a,b) = a + b - \sqrt{a^2 + b^2}$$

is locally Lipschitz on $\mathbb{R}^2$ and continuously differentiable away from $(0,0)$, but is not Clarke regular at the origin for this sign convention. For $(a,b) \neq (0,0)$,

$$\nabla\phi_c(a,b) = \left(1 - \frac{a}{\sqrt{a^2+b^2}}, 1 - \frac{b}{\sqrt{a^2+b^2}}\right).$$

At $(0,0)$,

$$\partial\phi_c(0,0) = \{(1,1) - \xi : \|\xi\| \leq 1\}.$$

Local Lipschitz continuity of the gradients of f, h, g, for example from C^2 data on compact tubes, together with joint local Lipschitz time dependence, makes E locally Lipschitz and gives a nonempty compact convex Clarke generalized Jacobian. Merely C^1 objective data do not suffice.

For an absolutely continuous trajectory $y(\cdot)$ that remains in such a tube, the composition $t \mapsto E(y(t),t)$ is absolutely continuous. The analysis requires the implemented selection to satisfy, almost everywhere,

$$\dot{E}(y(t),t) = G(t)\dot{y}(t) + E_t(y(t),t),$$

as stipulated in Assumption 5; existence of some chain-rule selection does not justify every Jacobian used in the dynamics. At an FB hinge, this compatibility and the spectral and drift bounds require a separate audit. For example, the directional derivative of ϕ_c at the origin along $(1,0)$ is zero, whereas its Clarke directional derivative is two.

C Nonconvex Branch-Lock Details

Definition 17 (Residual escape barrier and branch-lock margin)**.** Let $y_\star(t)$ be a continuous KKT branch with $E(y_\star(t),t) = 0$ on $[0,T_c]$. For a continuous radius schedule $r : [0,T_c] \to (0,\infty)$, define

$$\delta_b(t, r(t)) := \inf_{\|y - y_\star(t)\| = r(t)} \|E(y,t)\|. \tag{32}$$

If a residual envelope $\bar{s}(t)$ is available, the branch-lock margin is

$$M_i[r] := \inf_{t\in[0,T_c]} (\delta_b(t, r(t)) - \bar{s}(t)).$$

The radius schedule may be fixed, shrunk near a fold, or selected from a predeclared finite grid. Containment has branch-identity meaning only if the tube isolates the declared root. For $E(x) = x^3 - x$, a radius-three tube around -1 has boundary barrier six yet contains the distinct root $+1$. The constant trajectory at $+1$ has zero residual and never leaves that tube; a positive barrier alone does not identify the branch.

Theorem 18 (Moving-tube branch-lock certificate)**.** *Suppose E and $y_\star$ are continuous on the analyzed time window and y is continuous and $\|y(0) - y_\star(0)\| < r(0)$. If an audited residual envelope satisfies*

$$\|E(y(t),t)\| \leq \bar{s}(t) \quad \textit{and} \quad \bar{s}(t) < \delta_b(t, r(t)) \quad \textit{for all } t,$$

then $\|y(t) - y_\star(t)\| < r(t)$ for all $t \in [0,T_c]$.

Proof. The proof of Theorem 18 is a first-exit argument. If the trajectory first exits the moving tube at time τ, then $\|y(\tau) - y_\star(\tau)\| = r(\tau)$. By Definition 17, $\|E(y(\tau),\tau)\| \geq \delta_b(\tau, r(\tau))$, contradicting the audited envelope. □

Proposition 19 (A realizable nonconvex branch bottleneck)**.** *Consider the time-varying quartic*

$$f(x,t) = \frac{x^4}{4} - \frac{ax^2}{2} - b(t)x, \qquad a > 0,$$

with KKT residual $E(x,t) = x^3 - ax - b(t)$. For $|b(t)| < b_f := 2a^{3/2}/(3\sqrt{3})$, the residual has three real stationary branches: two local minima separated by a saddle. Along the left minimum branch $x_{\mathrm{L}}(t)$, the branch velocity satisfies

$$\dot{x}_{\mathrm{L}}(t) = \frac{\dot{b}(t)}{3x_{\mathrm{L}}(t)^2 - a}.$$

As $b(t) \uparrow b_f$, the left minimum and saddle coalesce, $3x_{\mathrm{L}}(t)^2 - a \to 0$, the branch gap shrinks to zero, and the escape barrier in (32) vanishes for any radius tied to that gap. Consequently, residual correctability and same-branch trackability separate on a smooth, realizable nonconvex optimization path.

Proof. The stationary equation is the depressed cubic $x^3 - ax - b = 0$, whose discriminant is $4a^3 - 27b^2$. Hence there are three real stationary points for $|b| < 2a^{3/2}/(3\sqrt{3})$, and the left minimum and middle saddle merge at the positive fold. Differentiating $E(x_{\mathrm{L}}(t), t) = 0$ gives $(3x_{\mathrm{L}}^2 - a)\dot{x}_{\mathrm{L}} - \dot{b} = 0$. The branch velocity diverges if $\dot{b}$ stays bounded away from zero as the Hessian vanishes; at the same time the minimum–saddle separation and any proportional escape barrier vanish. □

Proposition 20 (Constrained-manifold residual flatness)**.** *In dimension at least two, consider*

$$\min_{x:\, \|x\|^2 = 1} -c(t)^\top x,$$

with $c(t) \neq 0$ and KKT residual

$$E(x, \lambda, t) = \begin{bmatrix} -c(t) + 2\lambda x \\ \|x\|^2 - 1 \end{bmatrix}.$$

Let $a = \|c(t)\|$, $u = c(t)/a$, and $y_\star = (u, a/2)$. For every branch distance $r \in (0, 2)$, there exists a feasible point x_r with $\|x_r - u\| = r$ and multiplier $\lambda = a/2$ such that

$$\|E(x_r, \lambda, t)\| = ar.$$

Thus residual thresholds cannot uniformly identify the minimizer branch as the tangential curvature a vanishes.

Proof. Choose a unit tangent vector $v \perp u$ and angle θ satisfying $\|\cos\theta\, u + \sin\theta\, v - u\| = r$. Set $x_r = \cos\theta\, u + \sin\theta\, v$ and $\lambda = a/2$. Then $\|x_r\| = 1$, so the equality residual is zero, and the stationarity residual is $a(x_r - u)$, whose norm is ar. □

Proposition 21 (Quadratic-remainder escape-barrier bound)**.** *Fix a residual map $E(\cdot, t)$, a declared root $y_\star(t)$, and assume $E(y_\star(t), t) = 0$. Let $G_\star(t)$ be a fixed linearization at the root. Suppose that, for some $r > 0$ and $K \geq 0$,*

$$\|E(y_\star(t) + z, t) - E(y_\star(t), t) - G_\star(t)z\| \leq K\|z\|^2 \qquad \text{whenever } \|z\| \leq r.$$

Then the branch barrier at radius r satisfies

$$\delta_b(t, r) \geq \left[\sigma_{\min}(G_\star(t))r - Kr^2\right]_+. \tag{33}$$

The bound is valid for any residual geometry satisfying the stated local remainder condition; no sampling of the branch boundary is required.

Proof. For every z with $\|z\| = r$, the assumed remainder bound and the reverse triangle inequality give

$$\|E(y_\star(t) + z, t)\| \geq \|G_\star(t)z\| - K\|z\|^2 \geq \sigma_{\min}(G_\star(t))r - Kr^2.$$

Taking the infimum over the branch boundary and using nonnegativity of a residual norm proves (33). □

Example 22 (Unit-circle specialization)**.** For the unit-circle KKT residual in Proposition 20, write $z = (\Delta x, \Delta\lambda)$. The nonlinear remainder is

$$R(z) = \left(2\Delta\lambda\Delta x, \|\Delta x\|^2\right), \qquad \|R(z)\| \leq \sqrt{5}\,\|z\|^2.$$

Thus Proposition 21 applies with $K = \sqrt{5}$, recovering the circle-specific bound $\delta_b(t, r) \geq [\sigma_{\min}(G_\star(t))r - \sqrt{5}\,r^2]_+$.

Proposition 23 (Segment-wise digital residual enclosure)**.** *Let y_h linearly interpolate a proposed digital step from y_k to y_{k+1} on $I_k = [t_k, t_{k+1}]$. Partition I_k into intervals of length ℓ with midpoints m_j, and suppose $B_k \geq \sup_{t\in I_k} \left\| \frac{d^2}{dt^2} E(y_h(t), t) \right\|$. Then*

$$U_k := \max_j \left\{ \|E(y_h(m_j), m_j)\| + \frac{\ell}{2} \left\| \dot{E}(y_h(m_j), m_j) \right\| + \frac{B_k \ell^2}{8} \right\} \geq \sup_{t\in I_k} \|E(y_h(t), t)\| . \quad (34)$$

For the circle path, writing $v = (v_x, v_\lambda) = \dot{y}_h$, one may take

$$B_k = \sqrt{(a\omega^2 + 4|v_\lambda| \, \|v_x\|)^2 + 4 \, \|v_x\|^4}.$$

Thus $U_k < \inf_{t\in I_k} \underline{\delta}_b(t, r(t))$ is a sampling-independent commit test for the proposed segment.

Proof. Taylor's theorem about each midpoint bounds the residual vector by its value, first derivative, and the stated second-order remainder for every point in the subinterval. Taking norms and then the maximum proves (34). For the circle residual, direct differentiation along the linear segment gives $E'' = (-c'' + 4v_\lambda v_x,\ 2\|v_x\|^2)$, and $\|c''\| = a\omega^2$. □

Corollary 24 (Certified digital branch routing)**.** *If the initial digital state is inside the declared tube and every committed segment satisfies the commit test in Proposition 23, then its piecewise-linear digital path remains in the tube. If that tube contains no other KKT root, Algorithm 1 cannot emit* CertifiedIdentified *for a different branch under the audited barrier and arithmetic budgets.*

Proof. The segment enclosure supplies the residual envelope in Theorem 18; concatenate the finitely many committed segments. □

Proposition 25 (Verified re-anchor certificate)**.** *Let endpoint $k+1$ have arrived, let $F(y) = E(y, t_{k+1})$, and let $\widehat{y}$ be any tracker candidate corrected by damped Newton. Suppose a computable Jacobian inverse bound yields constants $\eta \geq 0$, $q_0 < 1$, and $L > 0$ such that the Newton–Kantorovich radius r satisfies*

$$r \geq \frac{\eta}{1 - q_0 - Lr}, \qquad q := q_0 + Lr < 1.$$

If the entire ball $B(\widehat{y}, r)$ satisfies the declared local-minimum condition (for the circle path, $\lambda \geq \lambda_{\min} > 0$), then the ball contains a unique KKT root. A router may therefore emit VerifiedReanchor *after the endpoint-arrival event; it makes no claim about anticipatory delivery or the identity of the path during the preceding abstention interval.*

Proof. The contraction form of the Newton–Kantorovich theorem gives existence and uniqueness in the ball when $q < 1$. The local-minimum condition excludes the nearby maximizer or saddle branch. The final statement follows because the certificate is endpoint-local. □

C.1 Verified active-set chart recovery

Let Θ be a certified uncertainty set for the current problem parameter. For an active set $A \subseteq \{1, \ldots, m\}$, replace the FB rows by the smooth chart equations

$$F_{A,\theta}(y) = \begin{bmatrix} \nabla_x L(x, \lambda, \mu; \theta) \\ h(x; \theta) \\ g_A(x; \theta) \\ \mu_{\bar{A}} \end{bmatrix} = 0. \quad (35)$$

The chart is semantically valid only if $\mu_A > 0$ and $-g_{\bar{A}} > 0$.

Theorem 26 (Robust active-set chart certificate)**.** *Let $\mathcal{C}$ be a finite set of candidate charts and $B_A = B(\widehat{y}_A, \rho_A)$. Suppose that, for every $\theta \in \Theta$, a computable corrector $Q_{A,\theta}$ maps B_A into itself, is a contraction there with a uniform factor $q_A < 1$, and its fixed points are exactly the roots of $F_{A,\theta}$. If the uniform sign margin*

$$\xi_A := \min \left\{ \inf_{\theta\in\Theta, y\in B_A, i\in A} \mu_i, \inf_{\theta\in\Theta, y\in B_A, j\notin A} -g_j(x; \theta) \right\} > 0, \quad (36)$$

then every $\theta \in \Theta$ has a unique root of chart A in B_A, and that root is an FB-KKT point with active set A. The router emits VerifiedReanchor *only if exactly one candidate chart passes; if zero or multiple charts pass, it abstains. Under the declared contraction, arithmetic, and parameter-uncertainty bounds, the delivered root therefore cannot be mislabeled with the wrong active set.*

Proof. Banach's fixed-point theorem gives a unique chart root in B_A, uniformly for each $\theta \in \Theta$. Equation (36) makes its active multipliers and inactive slacks strictly positive, so the chart equations imply the FB complementarity equations. Unique delivery is a router condition; when the evidence does not distinguish one chart, no active-set claim is emitted. This certifies the semantics of the delivered local root; without separate no-root certificates for the other charts, it does not assert that no KKT roots exist outside the audited balls. □

C.2 Exhaustive chart exclusion and dominance

The preceding result verifies the semantics of a returned local root, but does not identify it as the target when another local KKT chart remains possible. For the affine-chart case this gap can be closed exhaustively. Consider

$$\min_{x\in\mathbb{R}^n:\, Ax\le b} \ \tfrac{1}{2}x^\top Hx - \theta^\top x, \qquad \theta \in \Theta, \tag{37}$$

where H may be indefinite, Θ is compact, and the feasible polytope is nonempty and compact. For active set $\mathcal{A}$, its affine KKT chart has matrix

$$K_{\mathcal{A}} = \begin{bmatrix} H & A_{\mathcal{A}}^\top \\ A_{\mathcal{A}} & 0 \end{bmatrix}, \quad K_{\mathcal{A}} \begin{bmatrix} x_{\mathcal{A}}(\theta) \\ \mu_{\mathcal{A}}(\theta) \end{bmatrix} = \begin{bmatrix} \theta \\ b_{\mathcal{A}} \end{bmatrix}.$$

Let $v_{\mathcal{A}}(\theta)$ denote its objective value and let $Z_{\mathcal{A}}$ span $\ker A_{\mathcal{A}}$.

Theorem 27 (Exhaustive competitor-coverage certificate)**.** *Let $\mathcal{C}$ contain every active set that can represent a global minimizer of* (37)*; assume those minimizers satisfy LICQ and every consistent chart in $\mathcal{C}$ has nonsingular $K_{\mathcal{A}}$. Suppose a candidate $\mathcal{A}$ has, uniformly on Θ, strictly positive active multipliers and inactive slacks, and $Z_{\mathcal{A}}^\top H Z_{\mathcal{A}} \succ 0$. For every $\mathcal{B} \neq \mathcal{A}$, require at least one certified alternative:*

1. *$A_{\mathcal{B}}x = b_{\mathcal{B}}$ is inconsistent;*
2. *the affine chart root cannot satisfy primal–dual signs for any $\theta \in \Theta$;*
3. *$Z_{\mathcal{B}}^\top H Z_{\mathcal{B}}$ has a negative eigenvalue; or*
4. *the uniform dominance gap*

$$\Delta_{\mathcal{B}\leftarrow\mathcal{A}} := \inf_{\theta\in\Theta} \big(v_{\mathcal{B}}(\theta) - v_{\mathcal{A}}(\theta)\big) > 0.$$

Then $x_{\mathcal{A}}(\theta)$ is the unique global minimizer for every $\theta \in \Theta$. If any competitor lacks all four certificates, the router must abstain from target-identity delivery.

Proof. The candidate is a feasible strict local KKT minimum throughout Θ. Fix θ and suppose another global minimizer exists. By compactness it attains the minimum; by the coverage and LICQ assumptions it appears on a chart $\mathcal{B} \in \mathcal{C}$. A global minimizer cannot have inconsistent active equations, violate KKT signs, or admit a negative feasible-curvature direction. It must therefore fall under case 4, which gives $v_{\mathcal{B}}(\theta) > v_{\mathcal{A}}(\theta)$, a contradiction. Nonsingularity also precludes a second root on chart $\mathcal{A}$. □

The certificate is directly computable for a parameter box: chart roots and signs are affine, sign feasibility is a polytope-intersection test, and value differences are quadratic. In two parameter dimensions their exact box minimum is attained at a corner, an edge stationary point, or an interior stationary point. Higher-dimensional implementations may substitute a certified interval lower bound.

C.3 AMBIGUITY-ADAPTIVE NONLINEAR KRAWCZYK COVERAGE

Theorem 28 (Ambiguity-adaptive nonlinear chart coverage)**.** *Let a finite family of anchor tubes cover every KKT point that can attain the declared target over a parameter box Θ. On each tube, force every constraint with a strictly negative interval upper slack inactive and branch over the remaining k_r constraints. For delivery, require every resulting chart to be either (i) enclosed by a uniform Krawczyk inclusion with strict primal–dual sign and second-order margins, or (ii) rigorously excluded; any degenerate or unresolved leaf forces abstention. If exactly one enclosed chart uniformly dominates every other enclosed local minimum for the shared $\theta \in \Theta$, its target identity is sound. The search uses at most $\sum_r 2^{k_r}$ leaves and $\sum_r (2^{k_r+1} - 1)$ tree nodes, excluding per-leaf certificate cost.*

The proof follows the certificate construction below.

To establish the coverage and search bounds, let X_r be an anchor tube and Θ the current parameter box. If interval evaluation gives $\sup g_i(X_r, \Theta) < 0$, constraint i cannot be active in that tube and is removed. Let U_r contain the remaining k_r ambiguous inequalities. Every active set in the tube is then a subset of U_r, so a complete binary tree has at most 2^{k_r} leaves and $2^{k_r+1} - 1$ nodes.

For a leaf A, a root box Z_A, center z_A, real preconditioner C_A, and interval Jacobian $[J_A]$, define the uniform Krawczyk image

$$\mathcal{K}_A(Z_A, \Theta) = z_A - C_A F_A(z_A, \Theta) + (I - C_A[J_A](Z_A, \Theta))(Z_A - z_A). \tag{38}$$

The implementation requires $\mathcal{K}_A(Z_A, \Theta) \subset \operatorname{int} Z_A$ and an induced contraction bound $q_A < 1$. It then repeatedly intersects the root box with the Krawczyk image, without discarding any possible root. Active equalities whose interval range excludes zero, one-constraint charts whose interval tangent stationarity excludes zero, Krawczyk-disjoint boxes, infeasible primal–dual signs, and negative reduced curvature are certified exclusions. Zero sign or curvature margins are marked degenerate; failed inclusions and all merely numerical out-of-tube roots are unresolved.

Proof of Theorem 28. Fix $\theta \in \Theta$ and any target-attaining KKT point in an anchor tube. Strict interval slack pruning cannot remove one of its active constraints, so its active set occurs at a leaf of that tube's tree. A valid exclusion rules out that leaf as a local minimum. Otherwise (38), together with $q_A < 1$, gives one chart root in Z_A; strict sign margins make it an FB-KKT point and the second-order margin makes it a strict local minimum. Hence every target-attaining point covered by the anchor tubes is represented among the verified leaves. A strict interval upper bound on the shared-parameter difference $f_A(x_A, \theta) - f_B(x_B, \theta)$ eliminates each competitor B. Exactly one uniformly dominant leaf is therefore the unique target in the covered set. If the coverage premise contains every global minimizer, it is the unique global target. An unresolved or degenerate competitor destroys this implication, which is why the router abstains. The node bounds follow from the complete binary trees above. □

The proof deliberately separates the $O(m)$ interval-slack scan from chart count. Per-leaf interval algebra and sign checks still cost dimension-dependent work, so the theorem claims 2^{k_r}, not total constant time. Worst-case $k_r = m$ recovers exhaustive enumeration, and no statement is made about KKT roots outside the declared anchor tubes.

C.4 TARGET-AWARE GLOBAL ANCHOR COMPLETION

The preceding theorem is local to supplied tubes. We remove that premise on a declared compact domain by making coverage itself a verified output. Fix a parameter θ, include the domain boundary among the inequalities, and let $U = f(\bar{x}, \theta)$ be the objective value of any verified feasible incumbent $\bar{x}$.

Definition 29 (KKT-complete chart family)**.** For every $A \subseteq [m]$ with $|A| \le n$, let $H_A(z_A, \theta) = 0$ be a square chart on a compact search box D_A. The family is *KKT complete below U* if every KKT point x with active set A and $f(x, \theta) \le U$ is represented by a zero $z_A \in D_A$, from which its primal point and multipliers are interval recoverable. A primal–dual KKT chart with a certified multiplier box is one construction. In two dimensions, the reduced charts

$$\nabla f = 0, \qquad (g_i, v_i^\top \nabla f) = 0, \qquad (g_i, g_j) = 0, \quad v_i = (-\partial_2 g_i, \partial_1 g_i),$$

for $|A| = 0, 1, 2$, respectively, avoid an external multiplier bound; the multipliers are recovered from stationarity after root enclosure.

Theorem 30 (Closed-tree global anchor completion)**.** *Consider a C^2 inequality-constrained NLP on a compact domain, and suppose KKT necessity and LICQ hold at every global minimizer. Let a KKT-complete chart family below the verified incumbent value U be given. Subdivide every D_A with outward-rounded interval arithmetic. Require each leaf to be certified as one of:*

1. *target-irrelevant because $\inf[f](B) > U$;*
2. *root-free by an interval equation exclusion or a Krawczyk image disjoint from B;*
3. *a unique Krawczyk-enclosed root with strict primal–dual sign and second-order classification; or*
4. *unresolved or degenerate.*

No identity is delivered if any leaf has the fourth state. If all trees close and exactly one enclosed strict local minimum has a certified objective upper difference below zero against every other enclosed minimum, that root is the unique global minimizer. The chart count is at most $\sum_{j=0}^{n} \binom{m}{j}$, and the total interval work is $O(\sum_A N_A c_A)$, where N_A is the number of visited boxes and c_A is the per-box interval cost.

Proof. Let $x^\star$ be a global minimizer. Feasibility of $\bar{x}$ gives $f(x^\star, \theta) \le U$. By KKT necessity and LICQ, its active set has cardinality at most n, so KKT completeness places it at a zero of one enumerated chart. A leaf containing this zero cannot pass the objective test, because an inclusion function satisfies $\inf[f](B) \le f(x^\star, \theta) \le U$; nor can a valid equation or Krawczyk-disjoint test discard it. Closure therefore represents it in an enclosed root box. Strict sign and second-order tests remove roots that are infeasible, dual-infeasible, or not local minima. Hence every global minimizer is among the enclosed minima. A strict interval objective difference against every competitor proves that the sole winner has lower value than every other candidate, and therefore is the unique global minimizer. The chart and work bounds follow by summing the complete active-set family and its trees. □

Proposition 31 (Finite closure away from degeneracy)**.** *Assume additionally that every chart zero in the compact sublevel region $f \le U$ is isolated and nonsingular, distinct zeros have disjoint neighborhoods, every sign and second-order classification has a strict margin, and the interval extensions are convergent as box diameter vanishes. Require root-centered overlapping verification boxes near zeros, with the remaining domain still covered, so that subdivision boundaries cannot prevent strict inclusion. Then exhaustive longest-edge bisection with this verification step closes after finitely many subdivisions. If any nonsingularity, classification, or objective-dominance margin vanishes, finite delivery is not asserted and the correct status is unresolved.*

Proof. Each nonsingular zero has a neighborhood on which the Krawczyk inclusion is strict. On the compact complement of these neighborhoods, the chart residual is separated from zero; convergent interval extensions eventually expose a zero-free component on every sufficiently small box. Strict sign, curvature, and objective margins are likewise exposed after finite refinement. A finite subcover and the root-centered verification step yield finite closure. Bisection alone need not close: for $F(x) = x$, a root on a subdivision boundary never lies strictly inside either adjacent box. Vanishing margins remove precisely one of these finite-resolution arguments. □

Proposition 32 (Uniform parameter-box identity)**.** *Let Θ be a compact parameter box and suppose the hypotheses of Theorem 30 hold uniformly on Θ. In particular, every chart leaf is either excluded uniformly, or its Krawczyk image over $D_A \times \Theta$ is strictly contained in D_A with $q_A < 1$, and every sign, curvature, and pairwise objective-difference interval has a strict margin. Then the closed-tree router returns one* CertifiedGlobal *identity for every $\theta \in \Theta$. If the objective-difference interval contains zero for two enclosed roots, or a leaf is unresolved or degenerate, the router returns* Unresolved *for the whole box rather than selecting a pointwise winner.*

Proof. The uniform Krawczyk inclusion gives a unique root in each enclosed chart for every fixed $\theta \in \Theta$; the contraction bound is independent of the parameter choice. Uniform sign and curvature

intervals preserve KKT classification. The interval objective difference contains the value difference for every shared θ, so a strict negative upper bound orders the candidate against the competitor throughout Θ. Applying Theorem 30 pointwise therefore yields the same unique global target for the entire box. If any strict interval condition fails, the uniform implication is unavailable, and this strict contract requires abstention. □

C.5 Spectral minimax initializer

The proof of the spectral minimax initializer is elementary. Let $W \succeq 0$ be symmetric with spectrum contained in $[\underline{\lambda}, \bar{\lambda}]$, where $0 < \underline{\lambda} \le \bar{\lambda}$. For scalar gains $\gamma \ge 0$, the minimizer of

$$\max_{\lambda \in [\underline{\lambda}, \bar{\lambda}]} |1 - \gamma\lambda|$$

is

$$\gamma_s = \frac{2}{\underline{\lambda} + \bar{\lambda}}, \qquad \min_{\gamma \ge 0} \max_{\lambda \in [\underline{\lambda}, \bar{\lambda}]} |1 - \gamma\lambda| = \frac{\bar{\lambda} - \underline{\lambda}}{\bar{\lambda} + \underline{\lambda}}.$$

For any fixed $\gamma \ge 0$, the maximum of the convex function $\lambda \mapsto |1 - \gamma\lambda|$ over the interval is attained at an endpoint, so the objective is

$$m(\gamma) = \max\{|1 - \gamma\underline{\lambda}|,\ |1 - \gamma\bar{\lambda}|\}.$$

If $0 \le \gamma \le 1/\bar{\lambda}$, then $m(\gamma) = 1 - \gamma\underline{\lambda}$, which decreases with γ. If $\gamma \ge 1/\underline{\lambda}$, then $m(\gamma) = \gamma\bar{\lambda} - 1$, which increases with γ. Thus a minimizer lies in $[1/\bar{\lambda}, 1/\underline{\lambda}]$, where

$$m(\gamma) = \max\{1 - \gamma\underline{\lambda},\ \gamma\bar{\lambda} - 1\}.$$

The first term decreases and the second increases, so the unique minimizer equalizes them:

$$1 - \gamma\underline{\lambda} = \gamma\bar{\lambda} - 1,$$

which gives γ_s and the stated value.

Proposition 33 (Budget-aware gain preserves certificates under audited margins)**.** *Fix a residual tube $\mathcal{S}$ and implementation model. Let a budget-indexed implementation choose gains (γ_B, κ_B) from an audited candidate set Γ_B and report the corresponding margin $\rho_{\mathcal{S}}(\gamma_B, \kappa_B)$ from (18). If $\rho_{\mathcal{S}}(\gamma_B, \kappa_B) > 0$ and the corresponding certified-gain inequality holds, then Theorem 8 applies with these gains. Moreover, if γ_B is fixed and the implementation slack $R_e(s;\kappa)$ is independent of κ or is included in the audited value of $D_r(s;\kappa)$, then increasing κ cannot decrease the correction term in (17); certification is decided solely by the resulting audited margin.*

Proof. The first claim is exactly Theorem 8 applied to the selected pair (γ_B, κ_B). For the monotonicity statement, the correction part of (17) is

$$\gamma_B \underline{\lambda}\, \kappa (s^{2-q} + s^{2+q}),$$

which is affine and nondecreasing in κ for every $s > 0$. All other terms are either independent of κ under the stated condition or are already included in the audited $D_r(s;\kappa)$. Thus the certificate remains a direct audited-margin check, not a heuristic consequence of using a larger budget. □

C.6 Certified event tubes and extension boundaries

Public residual verification. For the separable QP, let the emitted allocation and equality multiplier be (x, ν), the arrived demand be d, capacities be u_i, offsets be c_i, and independent curvatures satisfy $q_i \in [\ell_i, h_i]$. Set

$$z_i(q_i) = x_i - \Pi_{[0,u_i]}(x_i - q_i x_i - c_i - \nu).$$

For $0 \le x_i \le u_i$, monotonicity of clipping makes z_i monotone in q_i. Consequently the exact worst fixed-candidate projected residual is

$$R^2 = \left(\sum_i x_i - d\right)^2 + \sum_i \max\{z_i(\ell_i)^2, z_i(h_i)^2\}.$$

Each coordinate attains its largest squared magnitude at an endpoint, and independence permits these endpoints to be chosen simultaneously. Thus two coordinatewise evaluations suffice; enumeration of 2^n corners is only a validation control. This certificate concerns the emitted primal–dual pair, not a reoptimized solution for each curvature. Numerical certification also requires an arithmetic error enclosure. The saved-binary64 exact audit above supplies that enclosure only for the archived values under its explicit data contract; it does not imply containment of a pre-rounded source-real input.

The original residual interface permits exact evaluation under the realized arrived-endpoint model. Its 186/470 deliveries at 5% uncertainty remain valid for that interface; they do not establish uniform residual validity over an unknown curvature box. Exhaustive corner checks retain 134/470 for the public interface, with corresponding counts 26/470 and 0/470 at 10% and 20%. The original point-event counts 47, 69, and 131 mixed endpoint errors with failure of a neighborhood-robustness requirement. Actual endpoint-label errors are 32, 57, and 119; at zero uncertainty they are zero, although 25 arrivals fail the neighborhood requirement.

On 200 new uniformly sampled demand pairs in $[0.08, 1.33]^2$, with the same QP and unchanged gates, public verification delivers 109, 18, 5, and zero at uncertainty levels 0, 5, 10, and 20%, respectively. Exact-model verification delivers 109, 37, 8, and zero. No public delivery has an endpoint-label error or violates the corner residual target; the closed-form bound agrees with corner enumeration on all 800 pairs. Separately, 100 new curvature boxes and all 64 corners give 76,800 successful event-enclosure checks. These constructed controls supplement, but do not add independent real-world trajectories to, the PJM replay.

For separable box-constrained quadratic layers, disjoint, ordered interval enclosures of lower/free and free/upper event times certify event membership and path order. Overlap leaves identity unresolved. Additive offset uncertainty requires wider endpoint intervals that account for the offset range.

For dense Hessians, Stieltjes/M-matrix structure preserves the monotone sensitivity argument and supports a global event enclosure. A generic SPD Hessian does not satisfy that ordering property; in that regime the interval–Krawczyk construction is only a local branch certificate and must abstain on boxes where contraction or sign margins fail.

On the scalar quartic branch, the a priori envelope scales inversely with fold clearance, requiring adaptive step control or abstention as the clearance vanishes. This argument applies to the scalar recurrence; convergence for general nonconvex problems remains outside its scope. External replay workloads test scope only and provide no evidence of live performance or alpha.

The supplementary eight-dimensional experiment evaluates an observed-only finite-pool heuristic. Its reference is best-of-192 local SLSQP, its labels record coarse sign and cap-activity patterns, and its iteration cap applies to each start separately. These empirical labels and per-start costs are distinct from the branch certificates and total CST budget analyzed here; observed label agreement does not establish nonlinear certification.