Research Square

Preprints are preliminary reports that have not undergone peer review.
They should not be considered conclusive, used to inform clinical practice,
or referenced by the media as validated information.

Certified Monotone-Face Krylov Acceleration for Box-Constrained Time-Fractional Boundary Control

Thiziri sifaoui

thiziri.sifaoui15@gmail.com

Department of Mathematics and Computer Science, Faculty of Sciences and Technology, University of Amine Elokkal El Hadj Moussa Eg Akhamouk, Tamanrasset, Algeri

KARIM Benalia

Department of Mathematical Sciences, Faculty of Hydrocarbons and Chemistry, University of Boumerdes, Boumerdes, Algeria

Sara bendjeddou

Department of Mathematics and Computer Science, Faculty of Sciences and Technology, University of Amine Elokkal El Hadj Moussa Eg Akhamouk, Tamanrasset, Algeri

---

Research Article

Keywords: numerical algorithms, time-fractional optimal control, discrete adjoint, active-set method, reduced conjugate gradients, projected Richardson iteration, spectral safeguard, box constraints

Posted Date: September 17th, 2026

DOI: https://doi.org/10.21203/rs.3.rs-10909868/v1

License: This work is licensed under a Creative Commons Attribution 4.0 International License.
Read Full License

Additional Declarations: No competing interests reported.

Certified Monotone-Face Krylov Acceleration for Box-Constrained Time-Fractional Boundary Control

Thiziri Sifaoui[1,2*] Karim Benalia[3†]

Sara Bendjeddou[2‡]

[1]LAROMAD Laboratory, Faculty of Sciences, Mouloud Mammeri University of Tizi Ouzou (UMMTO), Tizi Ouzou, Algeria

[2]Department of Mathematics and Computer Science, Faculty of Sciences and Technology, University of Amine Elokkal El Hadj Moussa Eg Akhamouk, Tamanrasset, Algeria

[3]Department of Mathematical Sciences, Faculty of Hydrocarbons and Chemistry, University of Boumerdes, Boumerdes, Algeria

Abstract

An L1 discretization of box-constrained time-fractional boundary control yields a strictly convex quadratic program whose reduced-Hessian actions are computed by transpose-consistent state and adjoint solves. We introduce a monotone-face principle and its active-set Krylov realization, SARC. A spectrally bounded projected-gradient step supplies a phase certificate; any feasible face correction that improves this safeguarded point inherits the same global R-linear objective bound. Under strict complementarity, a quantitative a posteriori neighborhood guarantees identification of the optimal face, after which truncated reduced CG terminates in at most the free-subspace dimension in exact arithmetic. An assembled row-sum bound gives a checkable spectral cap, while Lanczos provides a sharper matrix-free guide. On an 80-control verification problem, SARC reaches the direct discrete KKT solution in 10 phases and 41 Hessian products, compared with 158 projected iterations. Further tests examine control dimension, active-set accuracy, derivative consistency, fractional-order-dependent conditioning, and a heterogeneous L-shaped problem with 11,064 state unknowns.

Keywords: numerical algorithms; time-fractional optimal control; discrete adjoint; active-set method; reduced conjugate gradients; projected Richardson iteration; spectral safeguard; box constraints.

1 Introduction

Optimal control constrained by parabolic partial differential equations is a central model class in thermal regulation, diffusion management, process design, and distributed parameter systems; standard PDE-constrained optimization treatments are given by Hinze et al. (2009) and Tröltzsch (2010). Boundary actuation is particularly relevant when heaters, coolers, or flux-conditioning devices are accessible only at the boundary. When diffusion exhibits hereditary behavior, a Caputo time derivative provides a standard description of subdiffusive memory effects (Podlubny, 1999; Kilbas et al., 2006; Jin et al., 2019). Numerical approximation of such nonlocal dynamics has been studied through finite-difference, spectral, and L1-type schemes, including analyses for nonsmooth subdiffusion data (Lin and Xu, 2007; Jin et al., 2016). The resulting optimization problem combines nonlocal dynamics, boundary traces, and numerical linear algebra, and each of these ingredients affects the design of reliable solvers.

*Corresponding author. Email: thiziri.sifaoui15@gmail.com

†Email: k.benalia@univ-boumerdes.dz

‡Email: bendjeddousara@yahoo.com

Fractional optimal-control theory has developed along several complementary directions. Variational and Pontryagin-type formulations for Caputo systems are discussed by Frederico and Torres (2008), Bergounioux and Bourdin (2020), and Ndaïrou and Torres (2023), whereas fully discrete subdiffusion control with box constraints is analyzed by Jin et al. (2020), and Kien et al. (2022) establish optimality theory for fractional differential equations with control constraints. Fractional dynamics and boundary actuation are considered by Dorville et al. (2011) for a nonhomogeneous Dirichlet boundary problem, by Mehandiratta et al. (2021) for Caputo diffusion on metric graphs, and by Mehandiratta et al. (2026) for space–time fractional parabolic equations. Related fractional boundary-control theory on network geometries is developed by Leugering et al. (2023). For classical parabolic equations, finite-element analyses of Dirichlet boundary control include Gong et al. (2016), Liang et al. (2025), and the state-based treatment of Langer et al. (2025b). Building on these strands, the present work links a time-fractional Dirichlet-control discretization with reduced-Hessian spectral analysis, Picard relaxation, projected actuator constraints, Fermatean-informed penalty ambiguity, and fractional-order robustness. The transpose discrete adjoint keeps the boundary-control gradient consistent with the optimized finite-dimensional model.

A distinct source of uncertainty concerns the preference-based penalty coefficients in the objective. Fermatean fuzzy sets (Senapati and Yager, 2020) represent support, opposition, and hesitation within an enlarged membership–non-membership region; their principal extensions and decision-making applications are surveyed by Büyüközkan et al. (2024). In the present setting, the normalized score defines the central penalty, while a hesitation-dependent interval retains information lost through score-only defuzzification.

Uncertainty may also affect the fractional order. Risk-aware optimization with uncertain fractional exponents has been studied in fractional PDE control (Antil et al., 2021). Here we consider a finite set of plausible time-fractional orders and design one boundary-control policy that minimizes the worst scenario cost. This converts the memory-order uncertainty into a finite convex minimax program and exposes the trade-off between nominal and worst-case performance.

The computational analysis follows a *discretize-then-optimize* route. The continuous PDE defines the physical model, and the adjoint, reduced Hessian, fixed-point map, and convergence results are derived for its fully discrete counterpart. This treatment accommodates the boundary trace and fractional history directly in the assembled operators. Related state-based and nested-iteration approaches likewise emphasize the interaction between discretization, constraints, and efficient algebraic solvers (Langer et al., 2025a). The reduced structure identifies relaxed Picard as a preconditioned gradient method and yields a symmetric positive-definite formulation for conjugate gradients.

Projected Richardson is algebraically a preconditioned projected-gradient method, whereas PCG exploits the reduced SPD operator but does not by itself preserve bounds. Classical gradient-projection/CG and proportioning methods alternate face identification with conjugate-gradient exploration (Moré and Toraldo, 1991; Dostál, 1997, 2003; Dostál and Schöberl, 2005); primal–dual active-set methods provide a semismooth Newton interpretation (Hintermüller et al., 2002). SARC combines these ingredients through an operator-consistent acceptance mechanism for the L1 boundary-control reduced Hessian. The same forward–transpose-adjoint action supplies the control-metric safeguard and every free-face product. A scalar monotonicity test transfers the safeguard's objective contraction to a feasible face correction without imposing a contractivity condition on that correction. Bound crossings are reached by an explicit truncated step and trigger a restart. This construction yields a phase-wise complexity bound and a computable identification neighborhood while retaining a matrix-free implementation.

Table 1: Algorithmic position for box-constrained strictly convex quadratic programs.

Property	SARC	GPCG (Moré and Toraldo, 1991)	Proportioning (Dostál, 1997, 2003; Dostál and Schöberl, 2005)	PDAS (Hintermüller et al., 2002)
Feasible accepted iterates	Yes	Yes	Yes	Not intrinsic
Matrix-free face products	Yes	Possible	Possible	Possible
Face acceptance mechanism	Safeguard plus objective test	Projected/CG phase rules	Proportionality rules	Merit globalization
Convergence unit used here	Accepted phase	Outer iteration	Outer iteration	Newton step
Identification statement used here	Quantitative radius	Finite identification	Finite identification	Local active-set equivalence
Finite CG termination after identification	Yes	Yes	Method-dependent	Not CG-based

Algorithmic distinction from gradient-projection/CG frameworks

GPCG and proportioning methods established the effectiveness of alternating projected identification and subspace minimization (Moré and Toraldo, 1991; Dostál, 1997, 2003; Dostál and Schöberl, 2005). The present analysis isolates a modular interface between these operations. A projected safeguard establishes the phase estimate before the face routine is called, and the face routine is required to satisfy only feasibility and one scalar objective inequality. The projected-gradient rate used by the safeguard is classical; the contribution is the transfer of that rate to an otherwise unrestricted accepted face correction. The correction can therefore be produced by truncated CG, an exact principal solve, an inexact Krylov solve, or a semismooth step without changing the phase estimate. SARC is the reduced-CG realization of this interface, with explicit boundary truncation and restart rules.

To separate mathematical verification from geometric relevance, the experiments use two levels. The one-dimensional rod permits exact KKT comparison, detailed spectral diagnostics, and inexpensive parameter sweeps. A second benchmark then moves to a heterogeneous L-shaped plate with a re-entrant corner, four independently penalized boundary actuator patches, a spatially varying conductivity, and a protected interior zone receiving additional state weight. The same reduced-space algebra and projected iteration apply because they depend on the stacked matrices rather than on the spatial dimension.

Table 2 compares representative work on fractional control, boundary actuation, control bounds, and uncertainty. The present framework combines a fully discrete adjoint and reduced-Hessian structure with spectral analysis of the Picard feedback map, actuator constraints, hesitation-aware preference robustness, and minimax uncertainty in the time-fractional order.

Table 2: Representative positioning across fractional control, boundary actuation, constraints, and uncertainty.

Reference	Time-frac. dynamics	Boundary actuation	Box control	Uncertain frac. order	Fermatean preference
Dorville et al. (2011)	Yes	Yes	–	No	No
Gong et al. (2016)	No	Yes	Yes	No	No
Jin et al. (2020)	Yes	No	Yes	No	No
Mehandiratta et al. (2021)	Yes	Yes	–	No	No
Antil et al. (2021)	No[‡]	No	–	Yes	No
Kien et al. (2022)	Yes	No[†]	Yes	No	No
Leugering et al. (2023)	No[§]	Yes	–	No	No
Liang et al. (2025)	No	Yes	–	No	No
Langer et al. (2025b)	No	Yes	Yes	No	No
Mehandiratta et al. (2026)	Yes	Yes	–	No	No
Present work	Yes	Yes	Yes	Yes	Yes

"–" means that the feature is not a central component of the cited formulation. [†]Fractional differential-equation control rather than the present parabolic boundary-control PDE. [‡]Fractional spatial exponent in an elliptic diffusion model rather than a Caputo time derivative. [§]Fractionality enters the Sturm–Liouville spatial operator rather than through a Caputo time derivative.

The principal contribution is a monotone-face interface that transfers a certified projected-gradient objective rate to any feasible face correction satisfying one scalar acceptance test. Its SARC realization combines boundary truncation, matrix-free free-face CG, an explicit phase bound, a quantitative strict-complementarity radius, and exact-arithmetic finite termination after stable identification. A transpose-consistent L1 boundary-control discretization supplies the

application and the matrix-free operator. The numerical study examines certification, work counts, control scaling, derivative consistency, heterogeneous diffusion, and memory-order perturbations; preference and scenario uncertainty are used only to define secondary sensitivity tests.

The remainder of the paper is organized as follows. Section 2 introduces the Fermatean and fractional ingredients. Section 3 gives the one- and two-dimensional continuous and fully discrete models. Section 4 derives the exact discrete optimality system and reduced SPD structure. Section 5 develops spectral Picard, PCG, and projected control iterations. Section 6 presents the numerical study, including uncertainty and robustness experiments. Sections 7 and 8 discuss the findings and conclude.

2 Background and modeling conventions

2.1 Fermatean fuzzy numbers and score-based parameter elicitation

A Fermatean fuzzy number (FFN) is a pair

$$\widetilde{q} = (\mu, \nu), \qquad \mu, \nu \in [0,1], \qquad \mu^3 + \nu^3 \le 1,$$

where μ and ν are membership and non-membership degrees. The corresponding hesitation degree is

$$\pi = (1 - \mu^3 - \nu^3)^{1/3}.$$

We use the normalized score

$$S_N(\widetilde{q}) = \frac{1 + \mu^3 - \nu^3}{2} \in [0,1]. \tag{1}$$

Given a physically meaningful penalty interval $[q_{\min}, q_{\max}]$ with $0 < q_{\min} < q_{\max}$, the deterministic equivalent is

$$q^*(\widetilde{q}) = q_{\min} + S_N(\widetilde{q})(q_{\max} - q_{\min}). \tag{2}$$

Equation (2) maps the Fermatean assessment to the prescribed physical penalty range. Here μ represents support for a high penalty and ν represents opposition to that assessment. Equal values of $\mu^3 - \nu^3$ produce the same central penalty q^*.

2.2 Hesitation-aware ambiguity set

To retain information that is lost by the score, let $\pi = (1-\mu^3-\nu^3)^{1/3}$ and introduce an ambiguity level $\eta \in [0,1]$. Define

$$\delta(\widetilde{q}; \eta) = \eta\, \pi\, \min\{q^* - q_{\min}, q_{\max} - q^*\}, \tag{3}$$

and the penalty interval

$$\mathcal{Q}(\widetilde{q}; \eta) = [q^* - \delta(\widetilde{q}; \eta),\ q^* + \delta(\widetilde{q}; \eta)]. \tag{4}$$

The mapping has four useful properties: (i) $\mathcal{Q} \subseteq [q_{\min}, q_{\max}]$; (ii) $\eta = 0$ recovers the score-based model; (iii) $\pi = 0$ collapses the interval to q^*; and (iv), for fixed q^*, the radius is nondecreasing in η and π. Equal-score assessments with different hesitation degrees can therefore generate different ambiguity intervals.

Let $d(q^*) = \min\{q^* - q_{\min}, q_{\max} - q^*\}$. Within the separable family $\delta = \eta g(\pi) d(q^*)$, the baseline choice $g(\pi) = \pi$ satisfies normalization, continuity, and monotonicity. The alternative $g_r(\pi) = \pi^r$, $r > 0$, provides a sensitivity family: $r > 1$ narrows the interval and $0 < r < 1$ widens it. Monotonicity of the robust counterpart with respect to the upper penalties then preserves the corresponding ordering of robust values.

2.3 Caputo derivative

For $0 < \beta < 1$ and an absolutely continuous scalar function v, the left Caputo derivative is

$$ {}_0^C D_t^\beta v(t) = \frac{1}{\Gamma(1-\beta)} \int_0^t (t-s)^{-\beta} v'(s)\, ds. \tag{5} $$

At $\beta = 1$ it reduces to $v'(t)$. The nonlocal convolution in (5) encodes the full history of the state. We use the Caputo model at the physical level but derive optimization equations only after discretization.

3 Fermatean-informed time-fractional boundary-control model

3.1 Continuous physical model

Let $\Omega = (0, \ell)$ and $T > 0$. The temperature-like state $y(x,t)$ satisfies

$$ {}_0^C D_t^\beta y(x,t) = \partial_{xx} y(x,t), \qquad (x,t) \in \Omega \times (0,T), \tag{6} $$

with Dirichlet boundary controls

$$ y(0,t) = u_0(t), \qquad y(\ell,t) = u_\ell(t), \tag{7} $$

and prescribed initial state

$$ y(x,0) = y_0(x). \tag{8} $$

When actuator limits are present, the admissible controls additionally satisfy

$$ u_{\min} \le u_0(t), u_\ell(t) \le u_{\max} \quad \text{for a.e. } t \in (0,T), \tag{9} $$

with $u_{\min} < u_{\max}$. The unconstrained problem is recovered by omitting (9). The performance index is

$$ \begin{aligned} J(u_0, u_\ell) = {} & \frac{\alpha_f}{2} \int_0^\ell [y(x,T) - y_d(x)]^2\, dx + \frac{1}{2} \int_0^T \int_0^\ell y(x,t)^2\, dx\, dt \\ & + \frac{1}{2} \int_0^T \left[q_0^* u_0(t)^2 + q_\ell^* u_\ell(t)^2 \right] dt, \end{aligned} \tag{10} $$

where $q_0^*, q_\ell^* > 0$ are obtained from (2).

Remark 1 (Discrete analytical setting)**.** *The analysis is formulated for the discretization used in the computations. This yields a transpose-consistent adjoint for the Dirichlet boundary input and the L1 fractional-history operator. The results below therefore concern the finite-dimensional model defined in Sections 3–4.*

3.2 Spatial semi-discretization

Let M be the number of uniform spatial intervals, $h = \ell/M$, and $x_j = jh$, $j = 0, \dots, M$. We collect the interior values in

$$ z(t) = [y(x_1,t), \dots, y(x_{M-1},t)]^\top \in \mathbb{R}^{M-1}. $$

The centered second-order discretization gives

$$ {}_0^C D_t^\beta z(t) = A_h z(t) + B_h u(t), \qquad u(t) = \begin{bmatrix} u_0(t) \\ u_\ell(t) \end{bmatrix}, \tag{11} $$

where

$$A_h = \frac{1}{h^2}\begin{bmatrix} -2 & 1 & & \\ 1 & -2 & \ddots & \\ & \ddots & \ddots & 1 \\ & & 1 & -2 \end{bmatrix}, \qquad B_h = \frac{1}{h^2}\begin{bmatrix} e_1 & e_{M-1} \end{bmatrix}.$$

Thus the boundary controls enter the interior finite-difference equations through the first and last rows of the discrete Laplacian.

3.3 L1 time discretization and stacked state equation

Let N be the number of time steps, $\tau = T/N$, $t_n = n\tau$, and $z^n \approx z(t_n)$. For $0 < \beta < 1$ we use the L1 approximation

$$\delta_\tau^\beta z^n = \frac{\tau^{-\beta}}{\Gamma(2-\beta)} \sum_{j=0}^{n-1} a_j \left(z^{n-j} - z^{n-j-1} \right), \qquad a_j = (j+1)^{1-\beta} - j^{1-\beta}. \tag{12}$$

At $\beta = 1$ we use backward Euler,

$$\delta_\tau^1 z^n = \frac{z^n - z^{n-1}}{\tau}.$$

The fully discrete state equation is

$$\delta_\tau^\beta z^n = A_h z^n + B_h u^n, \qquad n = 1, \dots, N, \tag{13}$$

with z^0 sampled from y_0.

Stacking

$$Z = [(z^1)^\top, \dots, (z^N)^\top]^\top, \qquad U = [(u^1)^\top, \dots, (u^N)^\top]^\top,$$

we obtain the linear equality constraint

$$C_\beta Z - GU = d. \tag{14}$$

The matrix C_β is block lower triangular. Its diagonal blocks are $\gamma_\beta I - A_h$, where $\gamma_\beta > 0$ is the coefficient of z^n in (12) (or $1/\tau$ at $\beta = 1$). Because $-A_h$ is positive definite, every diagonal block is nonsingular and therefore C_β is nonsingular.

3.4 Discrete objective

Using rectangle quadrature in time and the standard interior-node quadrature in space, the discrete objective can be written

$$J_{h,\tau}(Z, U) = \frac{1}{2} Z^\top Q Z + c^\top Z + \frac{1}{2} U^\top R U + J_0, \tag{15}$$

where $Q \succ 0$ contains the running-state weights and the terminal penalty, $R \succ 0$ is block diagonal with blocks

$$\tau \ \mathrm{diag}(q_0^*, q_\ell^*),$$

c contains the terminal-target linear term, and J_0 is the constant part of the terminal tracking term. The numerical problem studied below is therefore

$$\min_{Z,U} \ J_{h,\tau}(Z, U) \qquad \text{subject to} \qquad C_\beta Z - GU = d. \tag{16}$$

3.5 Extension to a heterogeneous two-dimensional complex domain

The one-dimensional rod is used below for verification, but the algebraic construction is not dimension-specific. To test the method on a genuinely more structured geometry, we also consider the L-shaped domain

$$\Omega_2 = ((0, L_x) \times (0, L_y)) \setminus ([x_n, L_x] \times [y_n, L_y]), \tag{17}$$

with $L_x = 1.6$, $L_y = 1$, $x_n = 0.8$, and $y_n = 0.6$. The removed upper-right rectangle creates a re-entrant corner and two additional internal boundary segments. The state satisfies

$$ {}_0^C D_t^\beta y = \kappa_0 \nabla \cdot \big(k(x_1, x_2) \nabla y\big) \qquad \text{in } \Omega_2 \times (0, T), \tag{18}$$

where $\kappa_0 = 0.05$ is an overall diffusivity scale and $k(x_1, x_2) > 0$ describes material heterogeneity. Four disjoint boundary patches Γ_a, $a = 1, \ldots, 4$, are independently actuated,

$$y = u_a(t) \quad \text{on } \Gamma_a, \qquad y = 0 \quad \text{on } \partial\Omega_2 \setminus \bigcup_{a=1}^{4} \Gamma_a. \tag{19}$$

The four patches lie, respectively, on the left outer boundary, the lower outer boundary, the vertical notch wall, and the horizontal notch wall. Their locations are shown in Figure 10.

The two-dimensional objective preserves the quadratic structure while allowing a spatially weighted protected zone,

$$\begin{aligned} J_2(U) = {} & \frac{\alpha_f}{2} \int_{\Omega_2} w(x)[y(x, T) - y_d(x)]^2 \, dx + \frac{1}{2} \int_0^T \int_{\Omega_2} w(x) y(x, t)^2 \, dx \, dt \\ & + \frac{1}{2} \sum_{a=1}^{4} q_a^* \int_0^T u_a(t)^2 \, dt, \end{aligned} \tag{20}$$

where $w(x) \geq 1$ places additional emphasis near a protected component adjacent to the re-entrant corner. Each q_a^* is obtained from a Fermatean assessment using (2). In the benchmark the four assessments are $(0.85, 0.30)$, $(0.80, 0.35)$, $(0.78, 0.33)$, and $(0.88, 0.25)$.

A masked Cartesian grid is used on (17). The diffusion operator is assembled with a conservative five-point stencil and harmonic averaging of the conductivity across interior faces. At an actuated boundary face, the same face coefficient is transferred to the corresponding column of the boundary-input matrix. Non-actuated faces are held at the zero reference value. After applying the L1 time formula, the multidimensional problem again has the stacked form

$$C_\beta^{(2)} Z - G^{(2)} U = d^{(2)}, \tag{21}$$

with a positive diagonal control matrix $R^{(2)}$ and a positive state-weight matrix $Q^{(2)}$. Consequently, all reduced-Hessian, Picard, projection, and PCG results below apply to the two-dimensional discretization after replacing (C_β, G, Q, R) by $(C_\beta^{(2)}, G^{(2)}, Q^{(2)}, R^{(2)})$. This is the main reason for formulating the numerical theory at the fully discrete level.

4 Exact discrete optimality system

Proposition 1 (Existence and uniqueness)**.** *Let C_β be nonsingular, $Q \succ 0$, and $R \succ 0$. Then the fully discrete quadratic control problem has a unique optimal pair (Z^*, U^*). This statement applies to both the one-dimensional formulation* (16) *and the two-dimensional stacked model* (21)*.*

Proof. Since C_β is nonsingular, $Z = C_\beta^{-1}(d + GU)$. Substitution into (15) gives a reduced quadratic objective in U whose Hessian is

$$R + G^\top C_\beta^{-\top} Q C_\beta^{-1} G. \tag{22}$$

The first term is positive definite and the second is positive semidefinite. Hence the reduced Hessian is positive definite, so the reduced problem is strictly convex and coercive. It therefore has a unique minimizer U^*, and the corresponding Z^* is unique. □

Introduce a multiplier P for (14). The discrete Lagrangian is

$$\mathcal{L}(Z, U, P) = J_{h,\tau}(Z, U) + P^\top (C_\beta Z - GU - d).$$

Stationarity gives the exact finite-dimensional optimality system

$$QZ + c + C_\beta^\top P = 0, \tag{23a}$$
$$RU - G^\top P = 0, \tag{23b}$$
$$C_\beta Z - GU = d. \tag{23c}$$

Equivalently,

$$\begin{bmatrix} Q & 0 & C_\beta^\top \\ 0 & R & -G^\top \\ C_\beta & -G & 0 \end{bmatrix} \begin{bmatrix} Z \\ U \\ P \end{bmatrix} = \begin{bmatrix} -c \\ 0 \\ d \end{bmatrix}. \tag{24}$$

The one-dimensional reference solution in Section 6 is computed directly from (24). This provides an independent numerical target against which the Picard iteration is verified.

Remark 2 (Meaning of the discrete adjoint)**.** *The vector P is the adjoint of the discretized state map. The relation $C_\beta^\top P = -(QZ + c)$ incorporates the transpose of the L1 history operator and the discrete Dirichlet boundary input, producing the gradient of the finite-dimensional objective.*

4.1 Reduced Hessian and symmetric transformation

Eliminating the state gives a reduced quadratic objective $\phi(U)$ with Hessian

$$A_R = R + G^\top C_\beta^{-\top} Q C_\beta^{-1} G \succ 0.$$

The reduced first-order condition is $A_R U = f$ for a vector f determined by d and c. Let $V = R^{1/2} U$. Multiplication by $R^{-1/2}$ gives

$$(I + S)V = R^{-1/2} f, \qquad S = R^{-1/2} G^\top C_\beta^{-\top} Q C_\beta^{-1} G R^{-1/2} \succeq 0. \tag{25}$$

Thus the unconstrained reduced problem is symmetric positive definite. This observation permits a matrix-free PCG solve using forward and transpose state solves. PCG serves as the reference reduced-space iterative solver in the numerical study.

5 Spectral reduced-space algorithms

5.1 Affine fixed-point map

Equations (23a)–(23c) suggest a forward–adjoint–control sweep. Given a control U, define

$$Z(U) = C_\beta^{-1}(d + GU), \tag{26}$$
$$P(U) = -C_\beta^{-\top}[QZ(U) + c], \tag{27}$$
$$T(U) = R^{-1} G^\top P(U). \tag{28}$$

The optimal control satisfies $U^* = T(U^*)$. Substituting (26) into (28) gives

$$T(U) = b - HU, \tag{29}$$

with

$$H = R^{-1} G^\top C_\beta^{-\top} Q C_\beta^{-1} G, \tag{30}$$

$$b = -R^{-1} G^\top C_\beta^{-\top} (Q C_\beta^{-1} d + c). \tag{31}$$

Define the R-inner product and norm by

$$\langle v, w \rangle_R = v^\top R w, \qquad \|v\|_R = (v^\top R v)^{1/2}.$$

The matrix H is self-adjoint and positive semidefinite in this inner product because

$$R^{1/2} H R^{-1/2} = R^{-1/2} G^\top C_\beta^{-\top} Q C_\beta^{-1} G R^{-1/2} \succeq 0. \tag{32}$$

Hence all eigenvalues of H are real and nonnegative.

5.2 Relaxation interval

The relaxed Picard iteration is

$$U^{m+1} = U^m + \omega_m \big[T(U^m) - U^m \big]. \tag{33}$$

Let $e^m = U^m - U^*$. Using (29),

$$e^{m+1} = [I - \omega_m (I + H)]\, e^m. \tag{34}$$

Theorem 1 (Constant relaxation)**.** *Let $\lambda_{\max}$ be the largest eigenvalue of H. For a constant relaxation parameter $\omega > 0$, iteration* (33) *converges to the unique optimal control U^* for every initial control if and only if*

$$0 < \omega < \frac{2}{1 + \lambda_{\max}}. \tag{35}$$

Moreover, the asymptotic contraction factor in the R-norm is

$$\rho(\omega) = \max_i \left| 1 - \omega(1 + \lambda_i(H)) \right|. \tag{36}$$

Proof. By (32), H is diagonalizable in an R-orthonormal basis with eigenvalues $\lambda_i \geq 0$. In that basis, (34) multiplies the ith component by $1 - \omega(1 + \lambda_i)$. Convergence for all initial errors is equivalent to $|1 - \omega(1 + \lambda_i)| < 1$ for every i, which is exactly (35). Equation (36) follows immediately. □

Proposition 2 (Worst-case optimal stationary relaxation)**.** *Let $\lambda_{\min}$ and $\lambda_{\max}$ denote the extreme eigenvalues of H. Among all constant ω in the stable interval, the value minimizing the worst-case asymptotic factor* (36) *is*

$$\omega_\star = \frac{2}{2 + \lambda_{\min} + \lambda_{\max}}, \tag{37}$$

with

$$\rho_\star = \frac{\lambda_{\max} - \lambda_{\min}}{2 + \lambda_{\min} + \lambda_{\max}}. \tag{38}$$

Proof. The extremal amplification factors are $|1 - \omega(1 + \lambda_{\min})|$ and $|1 - \omega(1 + \lambda_{\max})|$. The minimax choice equalizes their magnitudes with opposite signs, which gives (37); substitution yields (38). □

The theorem explains why damping is not merely cosmetic. If $\lambda_{\max} > 1$, classical Picard ($\omega = 1$) is unstable. At the same time, a value of ω chosen too close to the upper boundary in (35) can converge very slowly because the dominant amplification factor approaches -1.

5.3 Certified spectral cap and Lanczos–secant acceleration

A single fixed relaxation requires prior tuning. Two implementations must be distinguished. If the reduced symmetric operator S in (25) is explicitly assembled, then

$$\lambda_{\max}(H) = \lambda_{\max}(S) \le \|S\|_2 \le \sqrt{\|S\|_1 \|S\|_\infty} = \|S\|_\infty =: \overline{\lambda}_{\max}, \tag{39}$$

because S is symmetric. The row-sum quantity in (39) is a directly checkable upper bound (up to outward rounding in floating-point arithmetic). Capping every step by $\theta/(1+\overline{\lambda}_{\max})$, with $0 < \theta < 2$, gives the certified variant covered by Propositions 3 and Theorem 2. For larger control spaces, forming S may be undesirable. The reported numerical variant instead uses a Lanczos estimate and residual backtracking; it is faster and matrix-free but is not advertised as a certified eigenvalue bound.

Define the fixed-point residual

$$r^m = T(U^m) - U^m. \tag{40}$$

For $m \ge 1$, the local secant estimate

$$\eta_m = \frac{\left\|T(U^m) - T(U^{m-1})\right\|_R}{\|U^m - U^{m-1}\|_R} \tag{41}$$

measures the observed feedback amplification along the latest displacement. The tentative relaxation is

$$\omega_m^{\text{trial}} = \min\left\{\frac{\theta}{1+\widehat{\lambda}_{\max}}, \frac{\widehat{\theta}}{1+\eta_m}, 1\right\}, \tag{42}$$

with $\theta = 1.90$ and $\widehat{\theta} = 1.80$ in the experiments. At the first iteration we start from 0.20 and apply the same numerical cap. The largest eigenvalue is estimated by Lanczos and denoted by $\widehat{\lambda}_{\max}$. Replacing $\widehat{\lambda}_{\max}$ by $\overline{\lambda}_{\max}$ from (39) converts the same rule into the certified variant.

For reproducibility, the residual backtracking is specified as follows. Starting from ω_m^{trial}, form a candidate with $\omega = \omega_m^{\text{trial}} 2^{-j}$, $j = 0, 1, \ldots$, and accept the first candidate satisfying

$$\|r(U^m + \omega r^m)\|_R \le \|r^m\|_R.$$

For the projected method, $U^m + \omega r^m$ is replaced by its R-projection and r by $\mathcal{R}_{\text{proj}}$. In the certified algorithm the accepted step is bounded below by a prescribed $\omega_{\min} > 0$ and above by $\theta/(1+\overline{\lambda}_{\max})$. Residual decrease is an acceptance diagnostic for the Lanczos-guided numerical variant, not a substitute for the hypotheses of Proposition 3. Accordingly, the theoretical and empirical variants are kept distinct in the statements and tables below.

Proposition 3 (Variable spectrally bounded relaxation)**.** *Suppose the accepted relaxation parameters satisfy*

$$0 < \omega_{\min} \le \omega_m \le \omega_{\max} < \frac{2}{1+\lambda_{\max}} \qquad \text{for all } m.$$

Then the sequence generated by (33) *converges linearly to* U^**.*

Proof. From (34), the ith spectral component is multiplied by $1 - \omega_m(1+\lambda_i)$. Continuity on the compact set $[\omega_{\min}, \omega_{\max}] \times [0, \lambda_{\max}]$ and strict inclusion in the stable interval imply the existence of $\kappa < 1$ such that

$$|1 - \omega_m(1+\lambda_i)| \le \kappa$$

for every m and i. Therefore $\|e^m\|_R \le \kappa^m \left\|e^0\right\|_R$. □

5.4 Box-constrained controls and projected Picard

Let

$$\mathcal{U}_{\mathrm{ad}} = [u_{\min}, u_{\max}]^{2N}.$$

Because R is diagonal, the orthogonal projection in the R-inner product is simply componentwise clipping. The constrained first-order condition is equivalent to the fixed-point relation

$$U^* = \Pi^R_{\mathcal{U}_{\mathrm{ad}}}\left(U^* + T(U^*) - U^*\right). \tag{43}$$

This motivates

$$U^{m+1} = \Pi^R_{\mathcal{U}_{\mathrm{ad}}}\left(U^m + \omega_m[T(U^m) - U^m]\right). \tag{44}$$

Theorem 2 (Global convergence with actuator bounds)**.** *For a constant ω satisfying $0 < \omega < 2/(1+\lambda_{\max})$, the projected iteration (44) converges globally and linearly to the unique minimizer of the box-constrained discrete quadratic program. The same conclusion holds for variable ω_m contained in a compact subinterval of this range.*

Proof. In the transformed variable $V = R^{1/2}U$, the unprojected gradient map has Jacobian $I - \omega(I+S)$ and therefore contraction factor $\max_i |1 - \omega(1+\lambda_i)| < 1$. Orthogonal projection onto a closed convex set is nonexpansive. The composition is thus a contraction. The fixed point is equivalent to the variational inequality for the strongly convex box-constrained quadratic objective, hence it is the unique optimizer. □

We measure constrained stationarity through

$$\mathcal{R}_{\mathrm{proj}}(U) = \Pi^R_{\mathcal{U}_{\mathrm{ad}}}(U + T(U) - U) - U.$$

The adaptive rule uses the same global spectral cap as in the unconstrained case and residual backtracking with $\|\mathcal{R}_{\mathrm{proj}}\|_R$.

5.5 Spectrally safeguarded active-set reduced CG

Projected Picard is globally reliable but discards conjugacy at every iteration. To retain feasibility while recovering Krylov acceleration, we introduce the spectrally safeguarded active-set reduced conjugate-gradient method (SARC). Define the reduced gradient $g(U) = A_R U - f$ and, with a tolerance $\epsilon_a \geq 0$, the binding estimate

$$\mathcal{A}(U) = \{i : U_i \leq u_{\min} + \epsilon_a,\ g_i(U) > 0\} \cup \{i : U_i \geq u_{\max} - \epsilon_a,\ g_i(U) < 0\}. \tag{45}$$

Its complement $\mathcal{F}(U)$ is the current free set. One SARC phase consists of:

Coordinates and metric. The analysis is expressed in the physical control U with the R-inner product. The implementation applies CG after the diagonal change of variables $V = R^{1/2}U$. Thus the free operator used by CG is the principal restriction of $B = R^{-1/2}A_R R^{-1/2}$, the bounds are scaled componentwise, and Euclidean CG products in V coincide with the R-metric formulation in U. Reported projected residuals are divided by $\sqrt{n_u}$. One Hessian product denotes one application of B, implemented by one incremental state solve and one transpose-adjoint solve; initial residual evaluations are counted separately whenever they require such an application.

1. a projected identification step $\widetilde{U} = \Pi^R_{\mathcal{U}_{\mathrm{ad}}}(U - \omega R^{-1} g(U))$, with $0 < \omega < 2/(1+\lambda_{\max})$;
2. construction of $\mathcal{A}(\widetilde{U})$ and the reduced free system

$$(A_R)_{\mathcal{F}\mathcal{F}} U_{\mathcal{F}} = f_{\mathcal{F}} - (A_R)_{\mathcal{F}\mathcal{A}} \widetilde{U}_{\mathcal{A}}; \tag{46}$$

3. CG iterations on (46). If a trial iterate crosses a bound, the Krylov step is truncated at the first encountered bound and the active set is rebuilt. A nondecreasing objective causes a restart from $\widetilde{U}$.

The reduced products in the CG phase are evaluated with one forward and one transpose solve; $(A_R)_{\mathcal{F}\mathcal{F}}p$ is obtained by zero-extending p outside $\mathcal{F}$. Thus neither A_R nor its free principal submatrix need be formed in a large-scale implementation. Here n_u denotes the number of discrete control variables.

For completeness, Algorithm 1 fixes the implementation details used below. The feasible step length is

$$\alpha_{\mathrm{box}} = \min\left\{\frac{u_{\max,i} - U_i}{p_i} : p_i > 0\right\} \wedge \min\left\{\frac{u_{\min,i} - U_i}{p_i} : p_i < 0\right\},$$

with an empty minimum interpreted as $+\infty$. A full CG step is accepted only when it is feasible and decreases the quadratic. Otherwise the phase returns its best feasible point and restarts. In the experiments, $\epsilon_a = 10^{-10}$, the projected-residual tolerance is 10^{-8}, and the inner iteration limit is the current free-set dimension.

Algorithm 1 (SARC phase and stopping rules). *1. Compute $g^k = A_R U^k - f$ and $\mathcal{R}_{\mathrm{proj}}(U^k)$. Stop if $\|\mathcal{R}_{\mathrm{proj}}(U^k)\|_R/\sqrt{n_u} \le \epsilon$.*

2. Set $\widetilde{U}^k = \Pi^R_{\mathcal{U}_{\mathrm{ad}}}(U^k - \omega_k R^{-1} g^k)$. For the objective-rate certificate use $0 < \omega_{\min} \le \omega_k \le 1/L$, where $L = \lambda_{\max}(R^{-1/2} A_R R^{-1/2})$; the wider interval $(0, 2/L)$ certifies convergence without the explicit objective-rate bound.

3. Form $\mathcal{A}(\widetilde{U}^k)$ from (45); fix these components. Set $z_0 = \widetilde{U}^k_{\mathcal{F}}$, $r_0 = f_{\mathcal{F}} - (A_R)_{\mathcal{F}\mathcal{A}}\widetilde{U}^k_{\mathcal{A}} - (A_R)_{\mathcal{F}\mathcal{F}} z_0$, and $p_0 = r_0$.

4. In scaled coordinates, compute $w_j = B_{\mathcal{F}\mathcal{F}} p_j$ and $\alpha_{\mathrm{cg}} = r_j^\top r_j / (p_j^\top w_j)$. If $\alpha_{\mathrm{cg}} > \alpha_{\mathrm{box}}$, truncate at the smallest nonnegative ratio; all components attaining that ratio are placed on their bounds, the point is clipped only at roundoff level, and the phase restarts. Otherwise apply the standard CG recurrences. Stop when $\|r_{j+1}\| \le \epsilon \max\{1, \|b_{\mathcal{F}} - B_{\mathcal{F}\mathcal{A}}\widetilde{V}_{\mathcal{A}}\|\}$.

5. Return the best feasible inner iterate if its objective does not exceed $\phi(\widetilde{U}^k)$; otherwise return $\widetilde{U}^k$. Rebuild the active set at the next phase.

Lemma 1 (Safeguard decrease). *Let $L = \lambda_{\max}(R^{-1/2} A_R R^{-1/2})$ and let $\widetilde{U} = \Pi^R_{\mathcal{U}_{\mathrm{ad}}}(U - \omega R^{-1} g(U))$ with $0 < \omega < 2/L$. Then*

$$\phi(\widetilde{U}) \le \phi(U) - \left(\frac{1}{\omega} - \frac{L}{2}\right)\|\widetilde{U} - U\|_R^2.$$

Proof. In scaled coordinates $V = R^{1/2}U$, projection optimality gives $\nabla\widehat{\phi}(V)^\top(\widetilde{V} - V) \le -\|\widetilde{V} - V\|^2/\omega$. Since the quadratic $\widehat{\phi}$ has Hessian $R^{-1/2}A_R R^{-1/2}$, the descent lemma is exact with curvature L. Combining the two relations and using $\|\widetilde{V} - V\| = \|\widetilde{U} - U\|_R$ proves the claim. □

Theorem 3 (Certified monotone-face acceleration). *Let $B = R^{-1/2}A_R R^{-1/2}$ have eigenvalues in $[\mu, L]$, $0 < \mu \le L$, and choose $0 < \omega_{\min} \le \omega_k \le \omega_{\max} \le 1/L$. At phase k, form the projected safeguard $\widetilde{U}^k$ and then allow* any *feasible face algorithm to return U^{k+1} satisfying $\phi(U^{k+1}) \le \phi(\widetilde{U}^k)$. Then*

$$\phi(U^k) - \phi(U^*) \le (1 - \mu\omega_{\min})^k[\phi(U^0) - \phi(U^*)]. \tag{47}$$

For $0 < \mu\omega_{\min} < 1$ and $0 < \varepsilon < \phi(U^0) - \phi(U^*)$, an absolute objective-gap tolerance ε is reached after at most

$$N_\varepsilon = \left\lceil \frac{\log((\phi(U^0) - \phi(U^*))/\varepsilon)}{-\log(1 - \mu\omega_{\min})} \right\rceil \tag{48}$$

phases. This conclusion applies to truncated CG, PDAS, semismooth Newton, or an inexact face solve without requiring the accepted correction itself to be contractive.

Proof. Use the scaled variable $V = R^{1/2}U$, let $C = R^{1/2}\mathcal{U}_{\mathrm{ad}}$, and write $F(V) = \widehat{\phi}(V) + \iota_C(V)$. The safeguard is $\widetilde{V}^k = \mathrm{proj}_C(V^k - \omega_k \nabla\widehat{\phi}(V^k))$. Projection optimality, applied with V^*, and L-smoothness give

$$F(\widetilde{V}^k) - F(V^*) \le \nabla\widehat{\phi}(V^k)^\top (V^k - V^*) - \frac{1}{2\omega_k}\|\widetilde{V}^k - V^k\|^2 - \frac{1}{2\omega_k}\|\widetilde{V}^k - V^*\|^2.$$

The projected-gradient three-point estimate, obtained by adding the projection variational inequality to the exact quadratic expansion of $\widehat{\phi}$, gives for every $W \in C$

$$F(\widetilde{V}^k) - F(W) \le \frac{1}{2\omega_k}\big(\|V^k - W\|^2 - \|\widetilde{V}^k - W\|^2\big) - \frac{\mu}{2}\|V^k - W\|^2.$$

The discarded remainder is $(1/\omega_k - L)\|\widetilde{V}^k - V^k\|^2/2 \ge 0$. Apply this estimate to the convex-combination comparison point $W = (1 - \mu\omega_k)V^k + \mu\omega_k V^* \in C$ and use the μ-strong convexity of $\widehat{\phi}$. Expansion and cancellation of the squared-distance terms yield

$$F(\widetilde{V}^k) - F(V^*) \le (1 - \mu\omega_k)[F(V^k) - F(V^*)].$$

On feasible points $F = \phi$. The monotone-face acceptance condition can only reduce the left-hand side. Since $\omega_k \ge \omega_{\min}$, iteration proves (47); solving that inequality for k gives (48). □

Remark 3 (Meaning of the identification constants)**.** *The three terms in* (49) *have separate roles. The first keeps every optimal free component at least ϵ_a from the bounds after the safeguard. The second preserves the strict-complementarity sign when the gradient is evaluated at the safeguarded point. The third makes the projected displacement large enough to place every optimal active component on its bound. Since σ, d, and the optimal partition are evaluated at a sufficiently accurate solution, ρ_{id} is an a posteriori diagnostic rather than a stopping threshold.*

Corollary 1 (Quantitative identification and finite termination)**.** *Assume strict complementarity. Let $\mathcal{A}^*$ and $\mathcal{F}^*$ be the optimal active and free sets, define*

$$\sigma = \min_{i \in \mathcal{A}^*} |g_i(U^*)|, \qquad d = \min_{i \in \mathcal{F}^*} \min\{U_i^* - u_{\min,i}, u_{\max,i} - U_i^*\},$$

and put $r_{\min} = \min_i R_{ii}$, $r_{\max} = \max_i R_{ii}$, $M_\infty = \|A_R\|_\infty$, and $K = 1 + \omega_{\max} M_\infty / r_{\min}$. If $0 \le \epsilon_a < d$ and

$$\|U^k - U^*\|_\infty < \rho_{\mathrm{id}} := \min\left\{\frac{d - \epsilon_a}{K}, \frac{\sigma}{2M_\infty K}, \frac{\omega_{\min}\sigma}{2r_{\max}}\right\}, \tag{49}$$

then the projected safeguard fixes every component of $\mathcal{A}^*$ at its correct bound and (45) *returns exactly $\mathcal{A}^*$. Once this neighborhood is entered, an unrestarted exact-arithmetic SARC face solve terminates in at most $|\mathcal{F}^*|$ CG steps.*

Proof. For a free component, nonexpansiveness of componentwise projection and $g_i(U^*) = 0$ give $|\widetilde{U}_i - U_i^*| \le K\|U - U^*\|_\infty$. Consider now $i \in \mathcal{A}^*$ at the lower bound. Strict complementarity gives $g_i(U^*) \ge \sigma$. The Lipschitz estimate $\|g(U) - g(U^*)\|_\infty \le M_\infty\|U - U^*\|_\infty$ implies $g_i(U) \ge \sigma/2$. Moreover, $0 \le U_i - u_{\min,i} < \rho_{\mathrm{id}}$, while $\omega_k g_i(U)/R_{ii} \ge \omega_{\min}\sigma/(2r_{\max}) \ge \rho_{\mathrm{id}}$. Hence the projected safeguard reaches $u_{\min,i}$. The second radius term preserves the positive gradient sign at the safeguarded point, so (45) selects i. The upper-bound case follows with reversed signs. If $i \in \mathcal{F}^*$, the first radius term keeps the safeguard farther than ϵ_a from either bound, and the component is not selected. Thus the estimated set is $\mathcal{A}^*$. Fixing these components gives the SPD principal system $(A_R)_{\mathcal{F}^*\mathcal{F}^*}U_{\mathcal{F}^*} = f_{\mathcal{F}^*} - (A_R)_{\mathcal{F}^*\mathcal{A}^*}U^*_{\mathcal{A}^*}$. Exact CG solves this system in at most $|\mathcal{F}^*|$ steps. Its solution is U^* and is strictly interior on $\mathcal{F}^*$, so a completed face solve introduces no new bound crossing. □

The safeguard supplies global convergence and the phase-complexity certificate independently of the selected face algorithm. SARC uses matrix-free CG, for which every free Hessian product is one forward–transpose-adjoint pair. Pure projected iteration converges throughout $(0, 2/L)$; the objective-rate estimate uses steps bounded by $1/L$. The Lanczos-guided experiments use the wider numerical rule stated above.

5.6 Robust counterparts

For the hesitation-aware penalty sets (4), consider

$$\min_{U} \max_{q_0 \in \mathcal{Q}_0,\, q_\ell \in \mathcal{Q}_\ell} \phi(U; q_0, q_\ell). \tag{50}$$

For every fixed U, the control-effort term is nondecreasing in q_0 and q_ℓ. Therefore the inner maximum is attained at the upper endpoints, and the robust counterpart is exactly the deterministic problem with $q_i = q_i^+$. This simple result is useful because it propagates hesitation without introducing an unnecessary inner numerical optimization.

For uncertain memory order, let $\mathcal{B} = \{\beta_1, \ldots, \beta_S\}$ and let $\phi_s(U)$ denote the reduced cost for scenario β_s. A single robust control is obtained from the convex epigraph program

$$\min_{U,\zeta} \ \zeta \quad \text{s.t.} \quad \phi_s(U) \le \zeta, \qquad s = 1, \ldots, S. \tag{51}$$

Every ϕ_s contains the same positive-definite control penalty, so $\max_s \phi_s(U)$ is strongly convex and the robust control is unique.

Proposition 4 (Scenario-dominance certificate)**.** *Let $U_{\bar{s}}^{\mathrm{or}}$ minimize one scenario cost $\phi_{\bar{s}}$. If*

$$\phi_s(U_{\bar{s}}^{\mathrm{or}}) \le \phi_{\bar{s}}(U_{\bar{s}}^{\mathrm{or}}) \qquad \textit{for every } s,$$

then $U_{\bar{s}}^{\mathrm{or}}$ is also the unique solution of the minimax problem (51)*, and scenario $\bar{s}$ is worst-case active at the robust optimum.*

Proof. For every control U, $\max_s \phi_s(U) \ge \phi_{\bar{s}}(U) \ge \phi_{\bar{s}}(U_{\bar{s}}^{\mathrm{or}})$. Under the stated cross-scenario inequalities, $\max_s \phi_s(U_{\bar{s}}^{\mathrm{or}}) = \phi_{\bar{s}}(U_{\bar{s}}^{\mathrm{or}})$, so the lower bound is attained. Strong convexity gives uniqueness. □

The certificate characterizes endpoint-dominated scenario grids, in which one scenario oracle also minimizes the maximum cost. The active-scenario pattern therefore follows directly from the cross-evaluation inequalities.

5.7 Implementation outline

The common assembly and spectral-safeguard steps are summarized below.

1. Input $\ell, T, \beta, M, N, \alpha_f$, the initial and target profiles, $[q_{\min}, q_{\max}]$, and the two Fermatean assessments.
2. Compute q_0^* and q_ℓ^* from (2).
3. Assemble $A_h, B_h, C_\beta, G, Q, R, c, d$.
4. Factorize C_β and $C_\beta^\top$ once.
5. Choose the mode: assemble S and compute $\overline{\lambda}_{\max} = \|S\|_\infty$ for the certified cap, or estimate $\widehat{\lambda}_{\max}$ by Lanczos for the matrix-free numerical guide.
6. Set $U^0 = 0$.
7. At iteration m, solve (26) and (27), compute $T(U^m)$ and r^m.
8. Stop if $\|r^m\|_R \le \varepsilon$.
9. Compute the secant estimate (41), choose the trial relaxation (42), enforce the certified cap when that mode is selected, and backtrack if the residual does not decrease.

10. For projected Picard, update $U^{m+1} = \Pi^R_{\mathcal{U}_{\mathrm{ad}}}(U^m + \omega_m r^m)$. For SARC, use this point to estimate (45), solve (46) by truncated CG, and restart whenever a bound is crossed.

The stopping criterion is therefore the fixed-point residual itself; the cost difference is reported only as a diagnostic and is not used as a termination test. The label "adaptive" in the numerical tables refers to the Lanczos–secant variant; the convergence theorems refer to steps satisfying their stated spectral bounds.

Table 3 clarifies iteration counts by recording the dominant work of each method. It prevents a misleading comparison between iterations that have different costs. Here one "state solve" means one application of C_β^{-1} and one "adjoint solve" one application of $C_\beta^{-\top}$; line-search trials may add reduced-gradient evaluations.

Table 3: Dominant work per outer iteration after factorization.

Method	State solves	Adjoint solves	Additional dominant work
Relaxed Picard	1	1	vector update and residual norm
Projected Picard	1	1	componentwise projection; extra solves if backtracking
Reduced PCG	1	1	inner products and recurrence vectors
SARC	1 per Hessian product	1 per Hessian product	active-set test and feasible restart
L-BFGS-B	at least 1	at least 1	limited-memory recursion and line search

6 Numerical experiments

6.1 Numerical settings

The benchmark uses a one-dimensional rod with

$$\ell = \pi, \qquad T = 1, \qquad y_0(x) = \sin x, \qquad y_d(x) = 0, \qquad \alpha_f = 1.$$

Unless stated otherwise, $M = 20$ spatial intervals, $N = 40$ time steps, $\beta = 0.8$, and residual tolerance 10^{-8}. The physical penalty interval is $[0.01, 0.10]$ and the Fermatean assessments

$$\widetilde{q}_0 = (0.85, 0.30), \qquad \widetilde{q}_\ell = (0.80, 0.35)$$

give $q_0^* = 0.081420625$ and $q_\ell^* = 0.076110625$. No random numbers enter the optimization problems. Python, NumPy, SciPy, pandas, and Matplotlib were used. Appendix B records the benchmark data, indexing, quadrature, actuator geometry, solver settings, and numerical sequences. The accompanying script assembles the 1D L1 system, reruns the new active-set comparison, writes residual histories, and records the software platform. The historical wall-clock measurements remain indicative because their original hardware metadata were not retained. No external observational dataset is used.

6.2 Spectral structure and unconstrained solver comparison

For the baseline problem,

$$\lambda_{\min}(H) = 0.354613, \qquad \lambda_{\max}(H) = 35.181365.$$

Hence constant Picard is stable only for

$$0 < \omega < 0.055277,$$

and Proposition 2 gives $\omega_\star = 0.053282$ with worst-case factor $\rho_\star = 0.927823$. Figure 1 shows the theoretical stationary contraction curve.

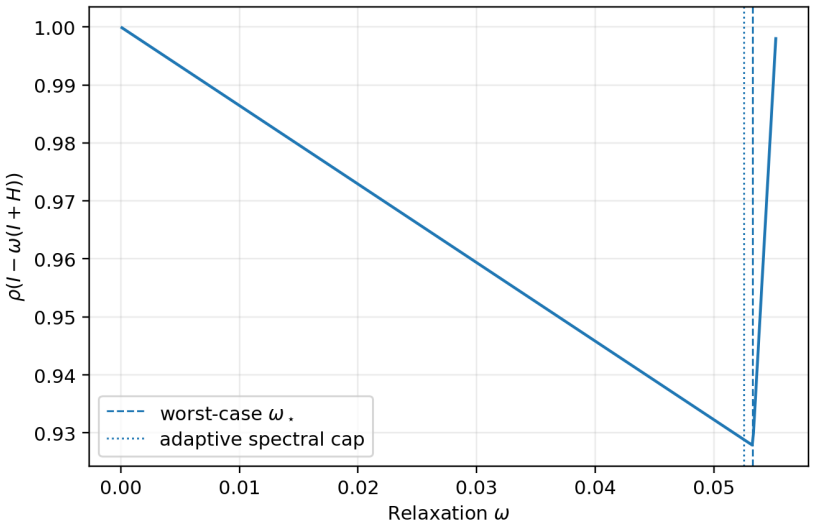


Figure 1: Worst-case spectral radius of stationary relaxed Picard.

Table 4 compares classical and relaxed Picard variants with reduced PCG and the discrete KKT solution. Classical Picard diverges because $\lambda_{\max} > 1$. The Lanczos–secant numerical variant reaches the discrete KKT solution in 185 iterations and closely matches the best tested fixed value. This row is an empirical acceleration result, whereas the certified cap is the variant covered directly by the spectral upper-bound argument. The worst-case-optimal stationary parameter minimizes the asymptotic spectral bound, although the iteration count also depends on the initial error. Reduced PCG exploits the SPD structure of (25) and reaches the same solution in 22 iterations.

Table 4: Baseline solver comparison ($\beta = 0.8$, $M = 20$, $N = 40$).

Method	Iterations	Final residual	Relative control error vs. KKT
Classical Picard	$6^\dagger$	2.14×10^{9}	2.66×10^{8}
Fixed Picard, $\omega = 0.03$	328	9.57×10^{-9}	2.95×10^{-8}
Near-best tested fixed, $\omega = 0.052$	186	9.65×10^{-9}	2.98×10^{-8}
Worst-case spectral $\omega_\star$	247	9.29×10^{-9}	1.15×10^{-9}
Lanczos–secant Picard	185	9.47×10^{-9}	2.90×10^{-8}
Reduced PCG	22	2.96×10^{-11}	1.54×10^{-10}
Exact KKT	–	–	–

†Stopped after the divergence threshold was exceeded. All converged methods give cost 0.1974321304.

Figure 2 shows the residual histories of the Picard variants. The fixed scan in Figure 3 confirms substantial sensitivity close to the stability boundary: the iteration count decreases to 186 at $\omega = 0.052$ but rises to 1830 at $\omega = 0.055$.

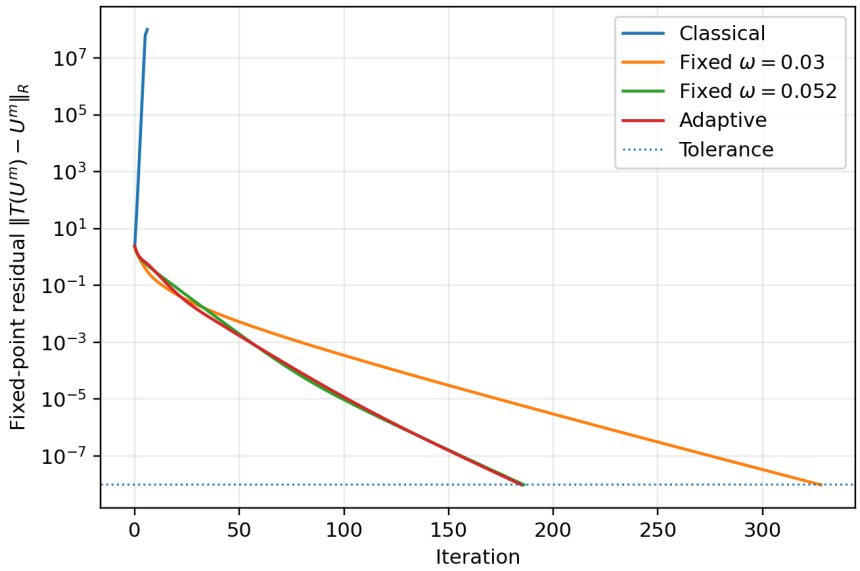


Figure 2: Fixed-point residual histories for the Picard variants.

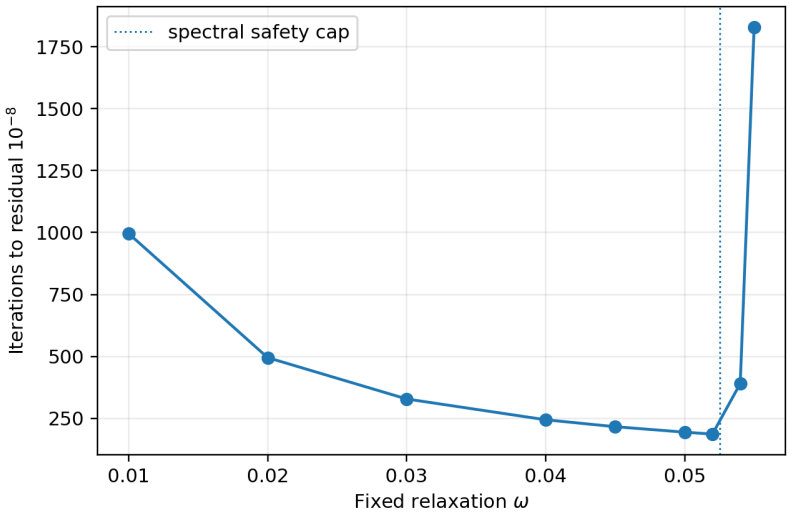


Figure 3: Iteration count under fixed relaxation; the vertical marker is the adaptive spectral cap.

6.3 Box-constrained actuators

The unconstrained optimal controls reach approximately -1.50 and -1.54 at the two boundaries, which motivates a bounded actuator test. We impose

$$-0.80 \le u_0^n, u_\ell^n \le 0.08.$$

The independent L-BFGS-B reference yields cost 0.20104148, and 18.75% of the control variables are active at one of the bounds. Lanczos–secant projected Picard reaches the same solution in 158 iterations, compared with 281 iterations for fixed $\omega = 0.03$; the relative control error against L-BFGS-B is about 2.0×10^{-7} in both cases. The comparisons based on reduced-Hessian products below provide a common matrix-free work unit.

Table 5: Box-constrained boundary control.

Method	Iterations	Cost	Active fraction
Projected Picard, $\omega = 0.03$	281	0.201041	0.1875
Projected Picard, Lanczos–secant	158	0.201041	0.1875
L-BFGS-B reference	21	0.201041	0.1875

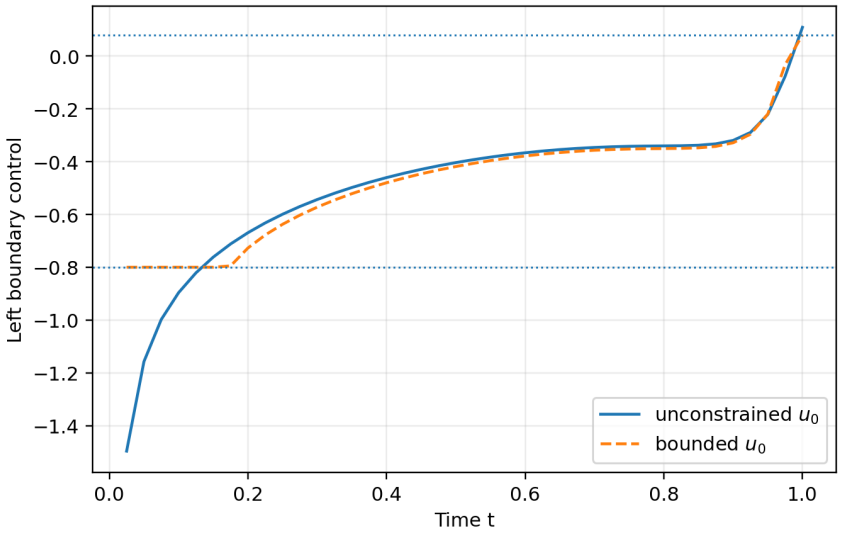


Figure 4: Effect of actuator bounds on the left boundary control.

This experiment changes the role of Picard: PCG is preferable for the unconstrained linear system, whereas projected Picard remains directly applicable when actuator bounds make the reduced optimality condition nonlinear through projection. L-BFGS-B provides an independently implemented reference solution.

6.4 Active-set Krylov acceleration

Table 6 evaluates the metric-consistent SARC implementation on the same constrained problem. A phase begins with a projected safeguard and may contain several free-subspace CG steps; reduced-Hessian products therefore provide the comparable matrix-free unit of work. Explicit truncation at the first encountered bound gives convergence in 10 phases and 41 Hessian products, with nine boundary restarts. The relative control error against L-BFGS-B is 1.15×10^{-7} and the scaled projected residual is 4.74×10^{-10}. The active fraction and objective coincide with both independent references. PDAS converges in four iterations, with an exact principal-system solve at each face update. Projected Picard uses 158 forward–adjoint sweeps.

Table 6: Work-normalized constrained comparison on the 1D benchmark.

Method	Phases/iter.	Hessian products	Restarts	Cost	Rel. control error
SARC	10	41	9	0.20104148	1.15×10^{-7}
PDAS	4	5[†]	–	0.20104148	1.16×10^{-7}
Projected Picard, Lanczos–secant	158	158	–	0.20104148	2.0×10^{-7}
L-BFGS-B	24	–	–	0.20104148	reference

[†]Gradient evaluations; PDAS additionally performs an exact dense solve on each free face, so this count is not a matrix-free work equivalence.

The accompanying script assembles the L1 matrices rather than reading precomputed objectives, applies the $R^{1/2}$ transformation used in the proofs, executes SARC, PDAS, and L-BFGS-B, and writes the complete residual history, comparison tables, and software-platform metadata. All methods use the same box and the same scaled projected-residual definition. The PDAS column counts gradient evaluations separately from dense principal-system solves; one SARC Hessian product represents one incremental state solve and one transpose-adjoint solve. No precomputed optimizer is read.

6.5 Spectral-bound audit and constrained scaling

For the baseline 80-control system, $\lambda_{\max}(H) = 35.181365$, whereas the certified row-sum bound is $\|S\|_\infty = 69.922395$. The bound is therefore 1.987 times the exact eigenvalue and changes the $\theta = 1.9$ cap from 0.052513 to 0.026790. This quantifies the price of certification: the row-sum cap is checkable but conservative, while Lanczos supplies the sharper matrix-free guide.

Table 7 evaluates the certified step $\omega = 0.99/L$. The factor $1 - \mu\omega = 0.962935$ gives a worst-case bound of 488 phases for a relative objective-gap reduction of 10^{-8}. The computed run reaches the PDAS solution in 12 phases. Evaluating the strict-complementarity quantities and the safeguard-displacement factor K at that solution gives the identification radius 2.64×10^{-6} in the scaled infinity norm.

Table 7: Audit of the certified monotone-face theorem on the 80-control problem.

μ	L	ω	Factor	Phase bound	Observed	ρ_{id}	Rel. error
1.3546	36.1814	0.027362	0.962935	488	12	2.64×10^{-6}	1.15×10^{-8}

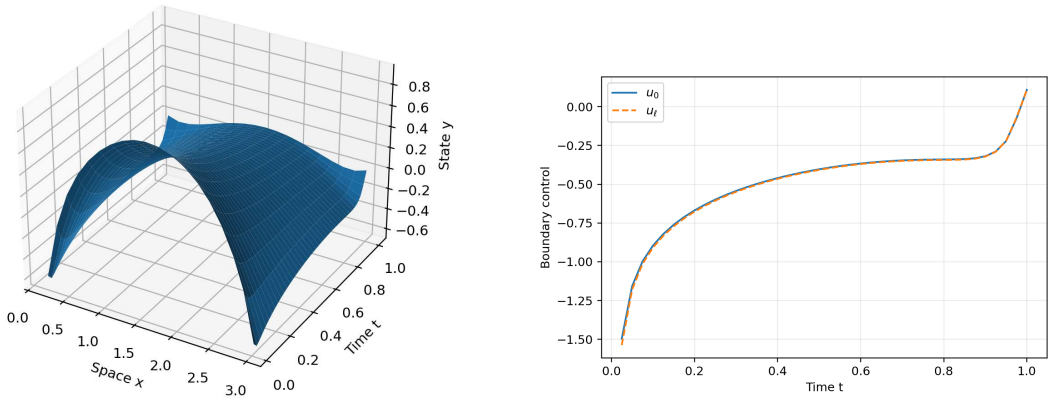
Table 8 varies the number of time-control coefficients without changing the physical horizon. SARC maintains the reference solution as the control dimension grows from 40 to 320. Its face phases use Hessian actions, whereas the reported PDAS implementation performs dense free-face solves.

Table 8: Constrained control-space scaling in the corrected R metric.

N	Controls	SARC phases	Hessian products	Restarts	PDAS iter.	SARC rel. error
20	40	6	28	5	3	1.29×10^{-7}
40	80	10	41	9	4	1.15×10^{-7}
80	160	13	50	12	4	4.02×10^{-8}
160	320	22	78	21	4	2.23×10^{-7}

6.6 Optimal state and controls

Figure 5 displays the discrete KKT state and unconstrained controls. Negative boundary values represent cooling action.



(a) Interior state.

(b) Boundary controls.

Figure 5: Baseline exact discrete solution.

6.7 Discrete operator and gradient verification

The reduced implementation was checked independently of the optimizer. For deterministic vectors p and q, the relative bilinear symmetry defect

$$\frac{|p^\top Bq - q^\top Bp|}{\max\{1, |p^\top Bq|, |q^\top Bp|\}}$$

is 8.75×10^{-16}. The eigenvalues of B lie in $[1.354613, 36.181365]$, confirming positive definiteness and giving condition number 26.71. A directional Taylor test compares the analytic reduced gradient $Bu - b$ with a forward difference of the objective. The first-order remainder decreases by a factor close to ten when the step is divided by ten, from 1.06×10^{1} at 10^{-1} to 1.04×10^{-6} at 10^{-8}. These tests exercise the transpose-consistent reduced action and objective implementation before any active-set decision. The vectors, seed, and complete sequence are generated by `reproduce_1d.py` and stored in machine-readable tables.

Table 9: Effect of the fractional order on the scaled reduced operator and SARC.

β	$\lambda_{\min}(B)$	$\lambda_{\max}(B)$	$\kappa(B)$	Phases	Hessian products	Restarts
0.6	1.7614	56.2649	31.94	12	45	11
0.7	1.5252	44.8793	29.43	11	43	10
0.8	1.3546	36.1814	26.71	10	41	9
0.9	1.2343	29.6500	24.02	9	40	8
1.0	1.1519	24.8676	21.59	8	36	7

The memory order changes more than the optimal objective: over this grid, decreasing β enlarges the upper spectral edge and increases the observed SARC work. This connects the fractional history operator directly to the solver behavior.

6.8 Fractional-order sensitivity

Table 10 varies β. The optimal cost increases over the tested range, while $\lambda_{\max}(H)$ decreases and the admissible Picard interval widens. These are benchmark-specific observations rather than general monotonicity claims.

Table 10: Fractional-order sensitivity.

β	KKT cost	Lanczos–secant iter.	$\lambda_{\max}(H)$
0.6	0.172747	223	55.2649
0.7	0.184263	204	43.8793
0.8	0.197432	185	35.1814
0.9	0.212247	167	28.6500
1.0	0.228592	156	23.8676

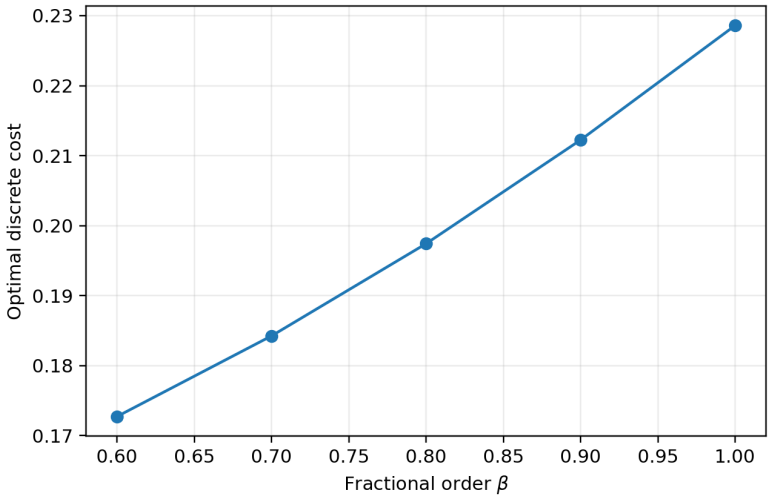


Figure 6: Optimal discrete cost versus fractional order.

6.9 Fermatean sensitivity and hesitation-aware robustness

Table 11 first reports score-based sensitivity. Cases S1 and S2 have the same normalized score, 0.7, but different hesitation. Under score-only defuzzification they both give $q_0^* = 0.073$ and exactly the same nominal cost, 0.196059.

Table 11: Fermatean-informed left-boundary penalty sensitivity.

Case	μ	ν	π	q_0^*	KKT cost
C1	0.7500	0.4000	0.8011	0.071104	0.195742
C2	0.8500	0.3000	0.7106	0.081421	0.197432
C3	0.9000	0.2000	0.6407	0.087445	0.198380
C4	0.9500	0.1500	0.5183	0.093430	0.199296
S1	0.7937	0.4642	0.7368	0.073000	0.196059
S2	0.8434	0.5848	0.5848	0.073000	0.196059

The ambiguity set (4) restores a role for hesitation. Table 12 varies the ambiguity level over $\eta \in \{0, 0.25, 0.5, 0.75, 1\}$. The robust control moves smoothly away from the score-only design as the ambiguity interval expands. At $\eta = 1$, the baseline FFNs generate upper penalties $q_0^+ = 0.094624$ and $q_\ell^+ = 0.094351$, and the robust optimum differs from the nominal control by 3.73% in the baseline R-norm. The improvement in worst-case objective remains modest because the uncertain quantities affect only the positive quadratic control-effort term; this is a calibration effect rather than a large performance gain.

Table 12: Hesitation-aware robustness as the ambiguity level η increases.

η	q_0^+	q_ℓ^+	Nominal-control worst case	Robust optimum	$\|U_{rob} - U_{nom}\|_R / \|U_{nom}\|_R$
0.00	0.081421	0.076111	0.197432	0.197432	0.0000
0.25	0.084721	0.080671	0.198710	0.198699	0.0098
0.50	0.088022	0.085231	0.199987	0.199945	0.0193
0.75	0.091323	0.089791	0.201265	0.201172	0.0284
1.00	0.094624	0.094351	0.202542	0.202379	0.0373

The equal-score cases make the role of hesitation especially transparent. Although S1 and S2 have the same score and hence the same nominal penalty $q_0^* = 0.073$, their different hesitation degrees produce different upper robust penalties and different robust objective values, as shown in Table 13. Thus the ambiguity construction changes the optimized decision problem rather than only relabeling the FFNs.

Table 13: Equal-score Fermatean assessments under full hesitation ambiguity ($\eta = 1$).

Case	Score	π	q_0^*	q_0^+	Robust optimal cost
S1	0.7000	0.7368	0.073000	0.092894	0.199215
S2	0.7000	0.5848	0.073000	0.088790	0.198588

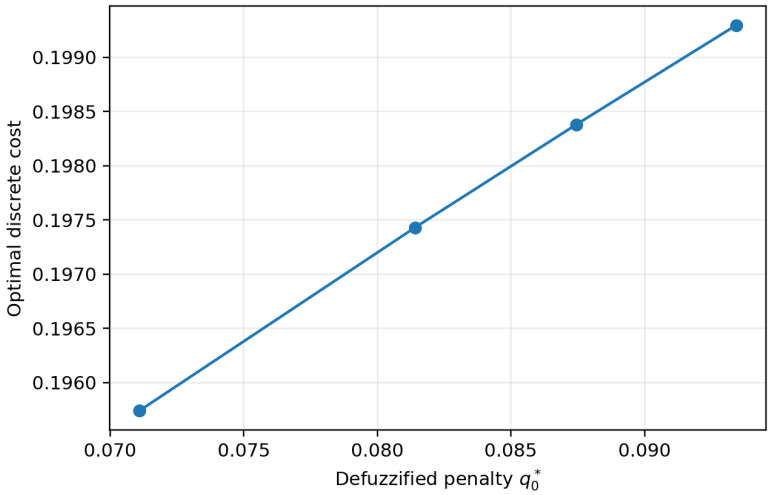


Figure 7: Nominal optimal cost versus score-based penalty for cases C1–C4.

6.10 Minimax robustness to the fractional order

The more consequential uncertainty experiment concerns the memory order itself. We use

$$\mathcal{B} = \{0.6, 0.7, 0.8, 0.9, 1.0\}.$$

This finite grid spans strongly subdiffusive behavior ($\beta = 0.6$), intermediate memory regimes, and the classical first-order limit ($\beta = 1$), with $\beta = 0.8$ used as the nominal design point. We compare the control optimized at $\beta = 0.8$ with the shared minimax control from (51). Table 14 reports the cross-scenario costs.

Table 14: Cross-scenario performance under fractional-order uncertainty.

β	Scenario oracle	Nominal $\beta = 0.8$ control	Minimax shared control	Active at minimax?
0.6	0.172747	0.182486	0.200008	No
0.7	0.184263	0.186446	0.198176	No
0.8	0.197432	0.197432	0.202927	No
0.9	0.212247	0.213915	0.213444	No
1.0	0.228592	0.234289	0.228592	Yes

The nominal control has worst-case cost 0.234289, whereas the minimax design reduces it to 0.228592, a 2.43% reduction. At the nominal order, its cost increases from 0.197432 to 0.202927 (2.78%). The test is endpoint dominated: the minimax value agrees with the $\beta = 1$ scenario-oracle value to the reported precision, and all other scenario costs remain below it. Proposition 4 identifies the $\beta = 1$ oracle as the minimax solution for this benchmark. Figure 8 shows the nominal and robust profiles across the scenario set.

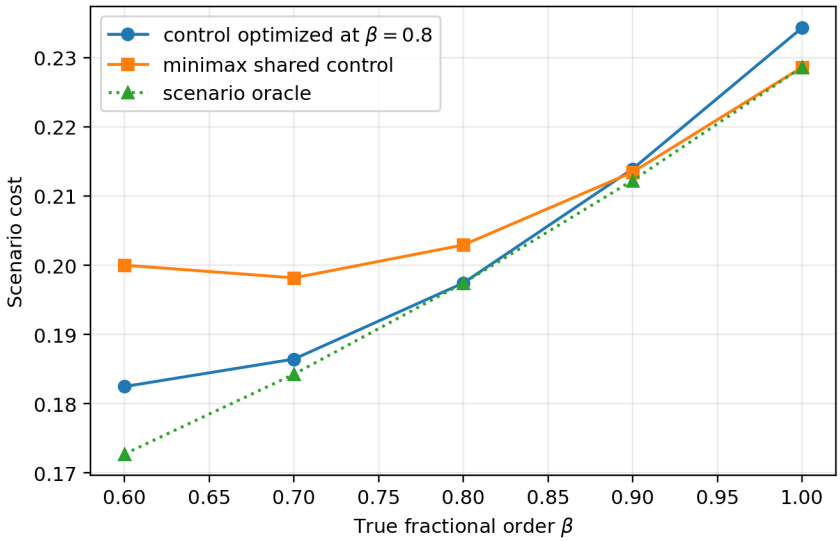


Figure 8: Nominal, minimax, and scenario-oracle costs under fractional-order misspecification.

6.11 Coupled mesh sensitivity and computational scaling

Table 15 reports the objective and its deviation from the finest available value. The KKT cost changes from 0.176738 on $(M, N) = (10, 20)$ to 0.211768 on $(60, 120)$, while the deviation decreases from 16.54% to 0.67% before the reference level. Because h and τ are refined simultaneously, these data establish numerical stabilization of the discrete objective but do not identify a spatial rate, a temporal L1 rate, or convergence of the discrete control to a continuous optimizer. Such claims would require separate refinement sequences, common-grid state and control errors, and regularity assumptions accounting for the initial-time singularity of subdiffusion. We therefore describe Table 15 only as a coupled mesh-sensitivity test.

Table 15: One-dimensional coupled refinement; the last column uses the finest available cost as a numerical reference.

M	N	KKT cost	Relative deviation (%)
10	20	0.1767381540	16.542
20	40	0.1974321304	6.770
30	60	0.2046302055	3.371
40	80	0.2082156492	1.678
50	100	0.2103511097	0.669
60	120	0.2117681090	0.000

For computational scaling, Table 16 increases the state dimension to 12,640 unknowns. PCG remains between 19 and 25 iterations, while Lanczos–secant Picard grows from 156 to 454 iterations at tolerance 10^{-7}. PCG is faster on these unconstrained problems; Picard supplies the forward–adjoint interpretation and extends directly to projected actuator bounds. Since hardware and environment metadata were not retained, the timing columns support only within-run qualitative comparison and not cross-platform performance claims.

Table 16: Matrix-free solver scaling ($N = 2M$); wall-clock times are environment specific.

M	N	State unknowns	Picard iter.	Picard time (s)	PCG iter.	PCG time (s)
20	40	760	156	0.050	19	0.003
40	80	3120	267	0.397	22	0.017
60	120	7080	365	3.882	25	0.120
80	160	12640	454	16.208	25	0.312

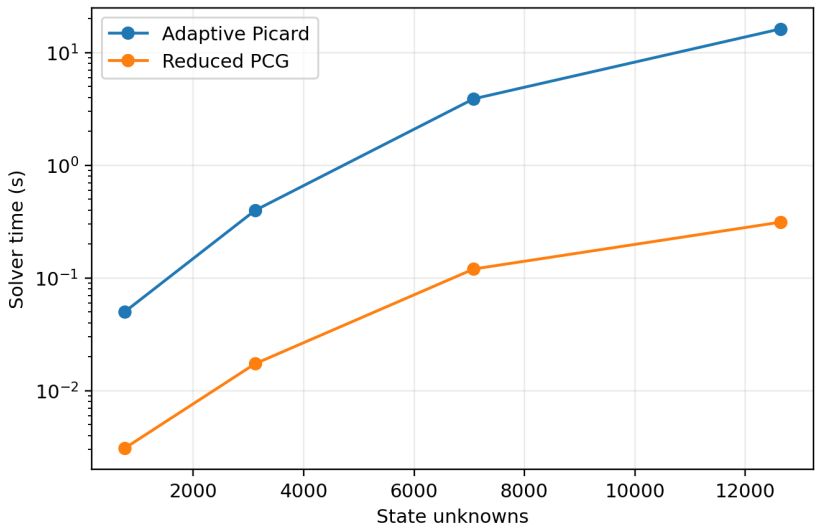


Figure 9: Indicative solver-time scaling in the supplied computational environment.

6.12 Heterogeneous two-dimensional L-shaped plate

The second benchmark is deliberately more demanding than the one-dimensional verification problem. It is an application-inspired, synthetic thermal-regulation problem on the L-shaped domain (17), with four distributed boundary actuators and spatially heterogeneous diffusion. Figure 10 shows the geometry, conductivity field, and actuator locations. The actuator patches

are

$$\Gamma_1 = \{0\} \times [0.18, 0.58], \quad \Gamma_2 = [0.84, 1.46] \times \{0\},$$
$$\Gamma_3 = \{0.8\} \times [0.68, 0.92], \quad \Gamma_4 = [0.94, 1.46] \times \{0.6\}.$$

The heterogeneous conductivity used in the numerical experiments is

$$k(x_1, x_2) = 0.78 + 0.38\,\mathbf{1}_{\{x_1<0.62\}} + 0.18\,\mathbf{1}_{\{x_2<0.30\}} + 0.26\exp\left[-\left(\frac{x_1 - 0.72}{0.22}\right)^2 - \left(\frac{x_2 - 0.43}{0.18}\right)^2\right]. \tag{52}$$

Two separated thermal loads define the initial state,

$$y_0 = 0.95\exp\left[-\left(\frac{x_1 - 0.34}{0.22}\right)^2 - \left(\frac{x_2 - 0.76}{0.18}\right)^2\right] + 1.25\exp\left[-\left(\frac{x_1 - 1.24}{0.25}\right)^2 - \left(\frac{x_2 - 0.27}{0.16}\right)^2\right] + 0.10\exp\left[-\left(\frac{x_1 - 0.68}{0.55}\right)^2 - \left(\frac{x_2 - 0.42}{0.42}\right)^2\right], \tag{53}$$

and the state weight is

$$w(x_1, x_2) = 1 + 3.2\exp\left[-\left(\frac{x_1 - 0.70}{0.18}\right)^2 - \left(\frac{x_2 - 0.48}{0.15}\right)^2\right]. \tag{54}$$

Thus the objective gives extra importance to the protected region near the re-entrant corner rather than treating every location as equally critical.

The baseline discretization uses 32 intervals in the x_1 direction, 20 intervals in x_2, and 24 time steps. The masked domain contains 461 spatial unknowns, producing 11,064 state unknowns and 96 control variables. The four Fermatean assessments stated in Section 3.5 give

$$(q_1^*, q_2^*, q_3^*, q_4^*) = (0.081421, 0.076111, 0.074738, 0.084963).$$

We impose the common actuator bounds

$$-0.20 \le u_a^n \le 0.05, \qquad a = 1, \dots, 4.$$

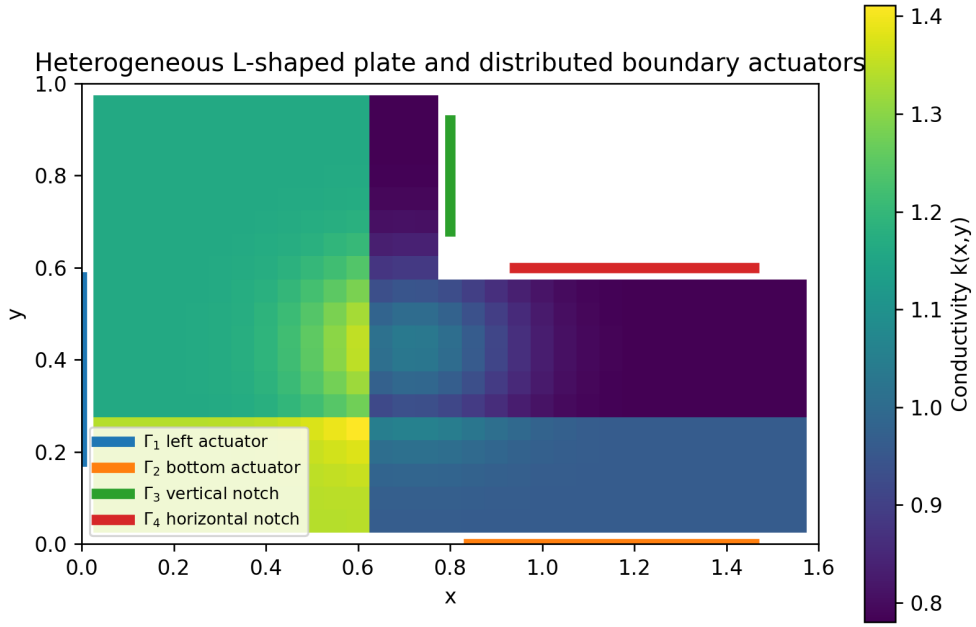


Figure 10: Heterogeneous L-shaped plate and the four distributed boundary actuator patches. The white upper-right region is removed from the physical domain.

Table 17 summarizes the main 2D results. The unconstrained PCG solution reduces the objective from 0.033003 to 0.025157 in 13 iterations. Under the actuator limits, L-BFGS-B gives cost 0.025249 with 14.58% of the discrete actuator values active at a bound. Projected Lanczos–secant Picard reaches the same constrained solution in 39 iterations; its relative control error with respect to L-BFGS-B is 1.95×10^{-6}. Relative to zero actuation, the constrained optimum reduces the total objective by 23.49%, the final maximum state from 0.212989 to 0.130814 (38.58%), and the protected-zone RMS state from 0.091431 to 0.069533 (23.95%).

Table 17: Heterogeneous L-shaped benchmark ($\beta = 0.8$).

Method	Iter.	Cost	Final RMS	Protected RMS	Active frac.
Uncontrolled	0	0.033003	0.092983	0.091431	0
Reduced PCG, unconstrained	13	0.025157	0.061186	0.069041	–
Projected Lanczos–secant Picard	39	0.025249	0.061596	0.069533	0.1458
L-BFGS-B constrained reference	16	0.025249	0.061596	0.069533	0.1458

Figure 11 shows that the two initial hot regions are progressively suppressed while respecting the nonconvex geometry and distributed actuator placement. Figure 12 shows the four control signals. The bottom and horizontal-notch actuators reach the cooling capacity over part of the horizon, so the box constraint is operationally active rather than cosmetic.

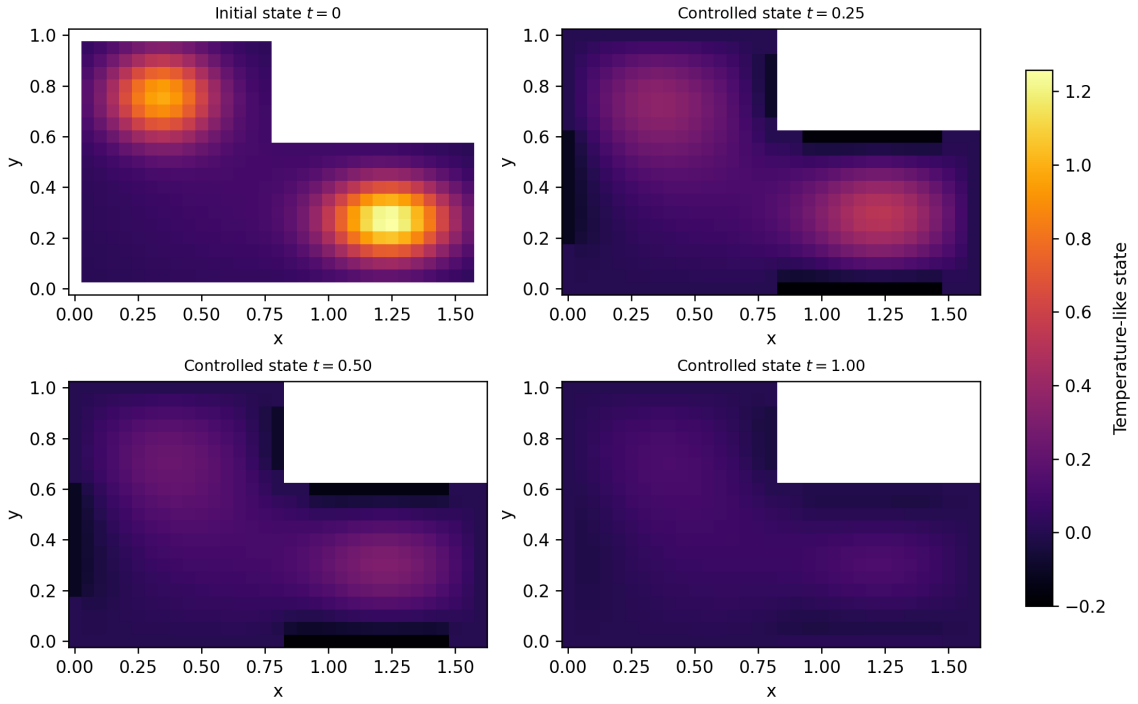
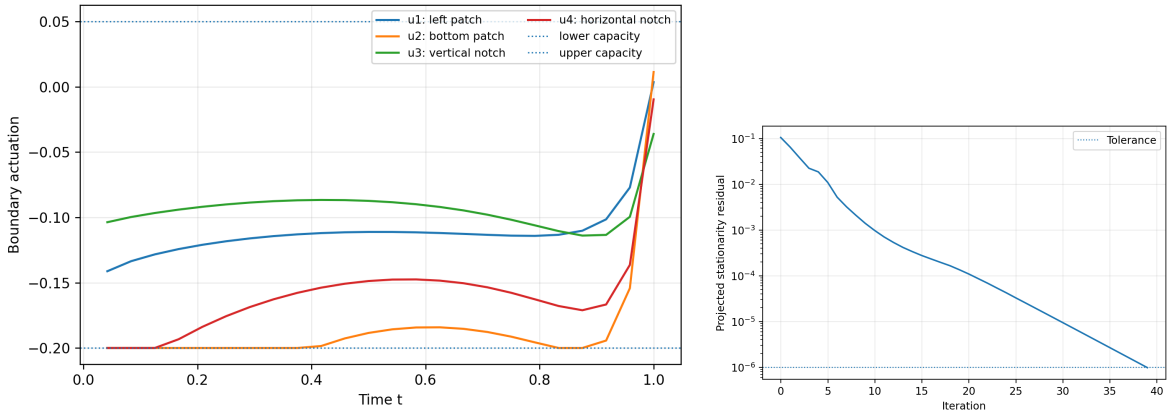


Figure 11: Initial and controlled state snapshots for the L-shaped heterogeneous plate.



(a) Four optimal actuator signals.

(b) Projected Picard residual.

Figure 12: Control saturation and convergence on the two-dimensional nonconvex-domain benchmark.

6.13 2D memory mismatch, material heterogeneity, and refinement

The constrained control designed at $\beta = 0.8$ is re-evaluated under $\beta = 0.6$ and $\beta = 1$. Table 18 compares the transferred policy with a scenario-specific constrained oracle. The largest relative regret is 1.77% at $\beta = 0.6$ and 1.35% at $\beta = 1$.

Table 18: Memory-order mismatch on the heterogeneous L-shaped plate.

β	Cost using $\beta = 0.8$ control	Scenario oracle	Relative regret
0.6	0.024853	0.024422	0.01766
0.8	0.025249	0.025249	0.00000
1.0	0.027475	0.027108	0.01353

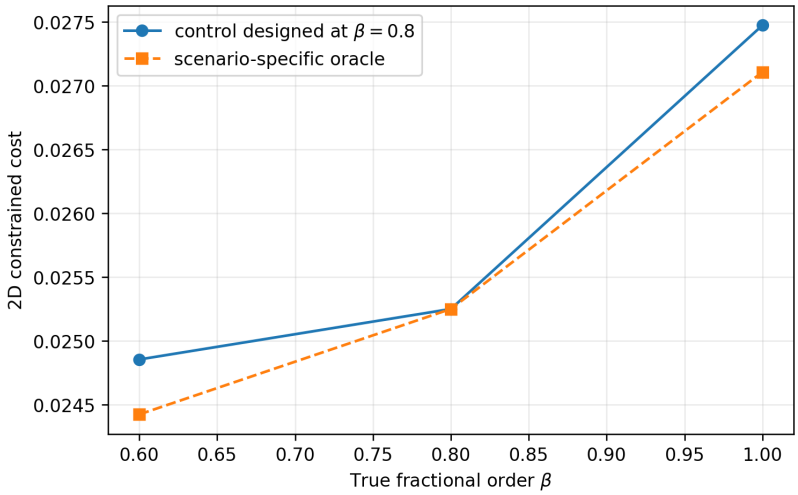


Figure 13: Cross-evaluation of the $\beta = 0.8$ constrained policy against scenario-specific 2D oracles.

A material ablation confirms that the variable conductivity is not numerically inert. On the same grid, the unconstrained PCG optimum has cost 0.024440 for homogeneous $k \equiv 1$ and 0.025157 for the heterogeneous field (52), an increase of approximately 2.93%. The final maximum state nevertheless falls from 0.142059 to 0.129205 because the heterogeneous field redistributes diffusion differently; this illustrates why cost, peak temperature, and protected-zone performance should be reported together.

Finally, Table 19 reports a three-level geometry-preserving mesh-sensitivity sequence. The control-space PCG iteration count remains between 11 and 13 while the stacked state dimension grows from 1,236 to 11,064. The successive cost changes decrease from 4.25% to 3.33%. The simultaneous space–time refinement and the re-entrant corner make this a mesh-sensitivity sequence rather than an estimate of asymptotic order.

Table 19: 2D geometry-preserving mesh-sensitivity sequence.

N_x	N_y	N_t	State unknowns	PCG iter.	Cost
16	10	12	1,236	11	0.023287
24	15	18	4,500	12	0.024320
32	20	24	11,064	13	0.025157

7 Discussion

Every optimization identity is exact for the finite-dimensional problem defined by the stated discretization, so the boundary gradient is consistent with the numerical state operator. Reduced PCG is substantially faster than Picard for the unconstrained quadratic problem. Under bounds, SARC combines projected identification with free-subspace CG: 41 reduced-Hessian products

replace 158 projected Picard sweeps while preserving feasibility and the reference solution. Theorem 3 applies to every feasible monotone face accelerator and supplies a worst-case phase bound. The certified run needs 12 phases against a bound of 488. The row-sum bound (39) gives a checkable cap when the reduced operator is assembled; its measured factor-of-1.987 conservatism quantifies the difference from the sharper Lanczos guide.

The operator audits complement optimizer agreement. Bilinear symmetry holds to machine precision, the directional Taylor remainder follows the expected first-order decay, and the positive spectral interval is measured directly. Across $\beta \in [0.6, 1]$, the scaled condition number decreases from 31.94 to 21.59 and SARC work decreases from 45 to 36 Hessian products. The fractional history therefore has a visible algebraic effect on both the reduced spectrum and the constrained iteration.

The box-constrained experiments yield nonzero active-set fractions in both spatial dimensions, and L-BFGS-B provides an independent numerical reference. SARC supplements projected identification with feasible free-subspace Krylov steps and a phase certificate that is independent of the accepted face routine. The Fermatean layer is a secondary penalty-elicitation sensitivity analysis rather than an additional optimization algorithm.

The one-dimensional minimax experiment is endpoint dominated, with the $\beta = 1$ oracle solving the finite-scenario problem. It verifies the dominance certificate but does not exercise a genuinely multi-active minimax optimum; the 2D mismatch test instead measures policy transfer across distinct memory regimes. The L-shaped plate demonstrates the same reduced-space structure with a re-entrant geometry, variable conductivity, four actuator patches, spatially nonuniform state weights, and more than 10^4 stacked state unknowns.

The multidimensional benchmark also changes how performance should be interpreted. In a heterogeneous medium, a lower peak state does not necessarily imply a lower quadratic objective because spatial redistribution, control effort, and protected-zone weighting compete. Reporting cost, maximum state, protected-zone RMS, active fraction, and cross-scenario regret therefore gives a more informative picture than a single convergence curve.

The numerical study uses linear synthetic benchmarks, a Cartesian representation of the L-shaped geometry, spatially uniform actuation on each patch, finite fractional-order scenario sets, and a uniform-grid L1 discretization. The convergence results concern the resulting finite-dimensional optimization problems, while the coupled refinement data quantify discretization sensitivity. The matrix-free formulation also accommodates graded temporal meshes, unstructured finite elements, larger actuator spaces, and multi-active scenario designs.

8 Conclusion

This paper developed a certified monotone-face acceleration framework for time-fractional boundary control. A single projected safeguard transfers its global R-linear objective rate to any feasible face solver that improves the safeguarded point. The result gives an explicit phase-complexity certificate and a computable strict-complementarity neighborhood for finite identification. SARC realizes the framework with truncated reduced CG, preserves feasibility, restarts safely at bound crossings, and inherits finite-dimensional CG termination after identification. The supporting spectral analysis characterizes Picard relaxation and supplies certified or Lanczos-guided safeguards.

The one-dimensional study isolates the mathematical structure: PCG reaches the discrete KKT solution in 22 iterations, while metric-consistent SARC matches the constrained L-BFGS-B reference using 10 phases and 41 Hessian products, versus 158 sweeps for adaptive projected Picard. PDAS supplies a second active-set reference, and tests with 40–320 controls retain close agreement with L-BFGS-B. Bilinear symmetry and directional Taylor tests verify the reduced action and gradient. Across the tested fractional orders, the SARC count follows the change in reduced condition number. The preference ambiguity set and finite-scenario minimax model

provide secondary sensitivity analyses; the latter reduces the tested worst-case cost by 2.43% relative to a policy optimized only at $\beta = 0.8$.

The two-dimensional L-shaped plate then tests whether these ideas survive a materially and geometrically richer problem. The baseline discretization contains 11,064 state unknowns and 96 controls. With four independently penalized boundary actuators and bounds $[-0.20, 0.05]$, the constrained optimum reduces the total objective by 23.49%, the final maximum state by 38.58%, and the protected-zone RMS state by 23.95% relative to zero actuation. Projected Lanczos–secant Picard agrees with the L-BFGS-B reference to a relative control error of 1.95×10^{-6}, and the transferred $\beta = 0.8$ policy incurs at most 1.77% regret over the tested $\beta \in \{0.6, 0.8, 1\}$ scenarios.

Compliance statements

- **Funding.** The authors received no funding for this work.
- **Conflict of interest.** The authors declare no conflict of interest.
- **Data and code availability.** All numerical data, model parameters, grid definitions, actuator specifications, solver tolerances, and results supporting the findings of this study are reported in the article and its appendices. The Python code used to generate the numerical experiments is available from the corresponding author upon reasonable request. No external observational data were used.
- **Author contributions.** T.S. contributed to conceptualization, methodology, mathematical analysis, software, numerical experiments, visualization, and preparation of the manuscript. K.B. contributed to problem formulation, mathematical analysis, and interpretation of the boundary-control model. S.B. contributed to numerical analysis, visualization, and editing. All authors reviewed and approved the final manuscript.

References

M. Hinze, R. Pinnau, M. Ulbrich, and S. Ulbrich. *Optimization with PDE Constraints*. Springer, Dordrecht, 2009.

F. Tröltzsch. *Optimal Control of Partial Differential Equations: Theory, Methods and Applications*. American Mathematical Society, Providence, 2010.

M. Bergounioux and L. Bourdin. Pontryagin maximum principle for general Caputo fractional optimal control problems with Bolza cost and terminal constraints. *ESAIM: Control, Optimisation and Calculus of Variations*, 26:35, 2020.

G. S. F. Frederico and D. F. M. Torres. Fractional optimal control in the sense of Caputo and the fractional Noether's theorem. *International Mathematical Forum*, 3(10):479–493, 2008.

W. Gong, M. Hinze, and Z. Zhou. A finite element method for Dirichlet boundary control problems governed by parabolic PDEs. *Journal of Scientific Computing*, 66(3):941–967, 2016.

B. Jin, R. Lazarov, and Z. Zhou. An analysis of the L1 scheme for the subdiffusion equation with nonsmooth data. *IMA Journal of Numerical Analysis*, 36(1):197–221, 2016.

B. Jin, R. Lazarov, and Z. Zhou. Numerical methods for time-fractional evolution equations with nonsmooth data: a concise overview. *Computational Methods in Applied Mathematics*, 19(1):33–57, 2019.

A. A. Kilbas, H. M. Srivastava, and J. J. Trujillo. *Theory and Applications of Fractional Differential Equations*. Elsevier, Amsterdam, 2006.

Y. Lin and C. Xu. Finite difference/spectral approximations for the time-fractional diffusion equation. *Journal of Computational Physics*, 225(2):1533–1552, 2007.

F. Ndaïrou and D. F. M. Torres. Pontryagin maximum principle for incommensurate fractional-orders optimal control problems. *Mathematics*, 11(19):4218, 2023.

H. Antil, D. P. Kouri, and J. Pfefferer. Risk-averse control of fractional diffusion with uncertain exponent. *SIAM Journal on Control and Optimization*, 59(2):1161–1187, 2021.

B. Jin, B. Li, and Z. Zhou. Pointwise-in-time error estimates for an optimal control problem with subdiffusion constraint. *IMA Journal of Numerical Analysis*, 40(1):377–404, 2020.

B. T. Kien, V. E. Fedorov, and T. D. Phuong. Optimal control problems governed by fractional differential equations with control constraints. *SIAM Journal on Control and Optimization*, 60(3):1732–1762, 2022.

I. Podlubny. *Fractional Differential Equations*. Academic Press, San Diego, 1999.

T. Senapati and R. R. Yager. Fermatean fuzzy sets. *Journal of Ambient Intelligence and Humanized Computing*, 11:663–674, 2020.

G. Büyüközkan, D. Uztürk, and Ö. Ilıcak. Fermatean fuzzy sets and its extensions: A systematic literature review. *Artificial Intelligence Review*, 57:138, 2024. doi:10.1007/s10462-024-10761-y.

R. Dorville, G. M. Mophou, and V. S. Valmorin. Optimal control of a nonhomogeneous Dirichlet boundary fractional diffusion equation. *Computers & Mathematics with Applications*, 62:1472–1481, 2011. doi:10.1016/j.camwa.2011.03.025.

U. Langer, R. Löscher, O. Steinbach, and H. Yang. State-based nested iteration solution of a class of optimal control problems with PDE constraints. *Mathematical Control and Related Fields*, 15(4):1496–1537, 2025. doi:10.3934/mcrf.2025043.

U. Langer, R. Löscher, O. Steinbach, and H. Yang. State-based approach to the numerical solution of Dirichlet boundary optimal control problems for the Laplace equation. *Applications of Mathematics*, 70:797–824, 2025. doi:10.21136/AM.2025.0166-25.

G. Leugering, G. Mophou, M. Moutamal, and M. Warma. Optimal control problems of parabolic fractional Sturm–Liouville equations in a star graph. *Mathematical Control and Related Fields*, 13(2):771–807, 2023. doi:10.3934/mcrf.2022015.

D. Liang, W. Gong, and X. Xie. A new error analysis for parabolic Dirichlet boundary control problems. *ESAIM: Mathematical Modelling and Numerical Analysis*, 59(2):749–787, 2025. doi:10.1051/m2an/2025006.

V. Mehandiratta, M. Mehra, and G. Leugering. Optimal control problems driven by time-fractional diffusion equations on metric graphs: Optimality system and finite difference approximation. *SIAM Journal on Control and Optimization*, 59(6):4216–4242, 2021. doi:10.1137/20M1340332.

V. Mehandiratta, M. Mehra, and G. Leugering. Optimal control problems governed by space–time fractional parabolic equations on metric graphs. *Evolution Equations and Control Theory*, 19:246–270, 2026. doi:10.3934/eect.2026016.

L. A. Zadeh. Fuzzy sets. *Information and Control*, 8(3):338–353, 1965.

J. Barzilai and J. M. Borwein. Two-point step size gradient methods. *IMA Journal of Numerical Analysis*, 8(1):141–148, 1988.

E. G. Birgin, J. M. Martínez, and M. Raydan. Nonmonotone spectral projected gradient methods on convex sets. *SIAM Journal on Optimization*, 10(4):1196–1211, 2000.

A. Beck and M. Teboulle. A fast iterative shrinkage-thresholding algorithm for linear inverse problems. *SIAM Journal on Imaging Sciences*, 2(1):183–202, 2009.

M. Hintermüller, K. Ito, and K. Kunisch. The primal-dual active set strategy as a semismooth Newton method. *SIAM Journal on Optimization*, 13(3):865–888, 2002.

J. J. Moré and G. Toraldo. On the solution of large quadratic programming problems with bound constraints. *SIAM Journal on Optimization*, 1(1):93–113, 1991. doi:10.1137/0801008.

Z. Dostál. Box constrained quadratic programming with proportioning and projections. *SIAM Journal on Optimization*, 7(3):871–887, 1997. doi:10.1137/S1052623494266250.

Z. Dostál. A proportioning based algorithm with rate of convergence for bound constrained quadratic programming. *Numerical Algorithms*, 34:293–302, 2003. doi:10.1023/B:NUMA.0000005347.98806.b2.

Z. Dostál and J. Schöberl. Minimizing quadratic functions subject to bound constraints with the rate of convergence and finite termination. *Computational Optimization and Applications*, 30:23–43, 2005. doi:10.1007/s10589-005-4557-7.

A Block form of the L1 operator

For completeness, the L1 derivative at time level n can be written

$$\delta_\tau^\beta z^n = \gamma_\beta a_0 z^n + \gamma_\beta \sum_{k=1}^{n-1} (a_{n-k} - a_{n-k-1}) z^k - \gamma_\beta a_{n-1} z^0,$$

where

$$\gamma_\beta = \frac{\tau^{-\beta}}{\Gamma(2-\beta)}.$$

This formula is used directly in the numerical implementation to assemble the lower-triangular time matrix. The right-hand side contribution associated with the prescribed initial state is $\gamma_\beta a_{n-1} z^0$.

B Complete benchmark specification

The 1D inputs are $\ell = \pi$, $T = 1$, $y_0(x) = \sin x$, $y_d = 0$, $\alpha_f = 1$, $(M, N, \beta) = (20, 40, 0.8)$, $[q_{\min}, q_{\max}] = [0.01, 0.10]$, and assessments $(0.85, 0.30)$ and $(0.80, 0.35)$. They give $(q_0^*, q_\ell^*) = (0.081420625, 0.076110625)$. The adaptive constants are $(\theta, \widehat{\theta}, \omega_0) = (1.90, 1.80, 0.20)$, the initial control is zero, the residual tolerance is 10^{-8}, the box test uses $[-0.80, 0.08]$, and the memory scenarios are $\{0.6, 0.7, 0.8, 0.9, 1.0\}$.

Spatial nodes are $x_j = j\ell/M$ and time levels are $t_n = nT/N$. States and controls are ordered first by time and then by spatial index or actuator number. Rectangle quadrature is used in time. The 1D spatial weight is h; in 2D each retained grid point carries its cell area multiplied by (54). The terminal block receives α_f, and control penalties use τ.

For 2D, $(L_x, L_y, x_n, y_n, \kappa_0, T) = (1.6, 1, 0.8, 0.6, 0.05, 1)$ and $(N_x, N_y, N_t) = (32, 20, 24)$. Equations (52)–(54) define conductivity, initial state, and weight. The four patches are $\Gamma_1 = \{0\} \times [0.18, 0.58]$, $\Gamma_2 = [0.84, 1.46] \times \{0\}$, $\Gamma_3 = \{0.8\} \times [0.68, 0.92]$, and $\Gamma_4 = [0.94, 1.46] \times \{0.6\}$. Their assessments are $(0.85, 0.30)$, $(0.80, 0.35)$, $(0.78, 0.33)$, and $(0.88, 0.25)$; their penalties are

$(0.081421, 0.076111, 0.074738, 0.084963)$ and the control interval is $[-0.20, 0.05]$. Interior-face conductivities are harmonic means; actuated boundary-face coefficients enter the corresponding input columns and other boundary values are zero.

The fixed-relaxation pairs $(\omega, \text{iterations})$ are $(0.01, 996)$, $(0.02, 495)$, $(0.03, 328)$, $(0.04, 244)$, $(0.045, 216)$, $(0.05, 194)$, $(0.052, 186)$, $(0.054, 390)$, and $(0.055, 1830)$; each converges below 10^{-8} to cost 0.1974321304. The 1D refinement pairs (M, N) are $(10, 20)$, $(20, 40)$, $(30, 60)$, $(40, 80)$, $(50, 100)$, and $(60, 120)$, with KKT costs 0.1767381540, 0.1974321304, 0.2046302055, 0.2082156492, 0.2103511097, and 0.2117681090. The 2D refinement triples are $(16, 10, 12)$, $(24, 15, 18)$, and $(32, 20, 24)$, with 103, 250, and 461 spatial unknowns. The accompanying CSV tables contain these numerical sequences in machine-readable form.

The program `code/reproduce_1d.py` assembles A_h, B_h, the L1 history matrix, Q, R, d, and the reduced Hessian. It performs the $R^{1/2}$ scaling, runs SARC, PDAS, and L-BFGS-B, and writes the solver comparison, residual history, spectral audit, scaling study, certificate audit, and software environment to `results/`. No precomputed optimizer or objective value is used as input.