Mathematics

https://www.aimspress.com/journal/Math

AIMS Mathematics, 11(7): 21385–21411.
DOI: 10.3934/math.2026867
Received: 17 March 2026
Revised: 28 May 2026
Accepted: 24 June 2026
Published: 17 July 2026

Research article

Geometric enforcing of sequential convergence in fixed-time stability: An anisotropic vector field approach

Xiaotian Liang[1,2], Yanbo Wang[3], Lei Yang[3], Hui Cao[1,2] and Shuangsi Xue[1,2,*]

[1] National Key Laboratory for High Energy Pulsed Power, Xi'an Jiaotong University, Xi'an 710048, China

[2] School of Electrical Engineering, Xi'an Jiaotong University, Xi'an 710049, China

[3] School of Automation and Information Engineering, Xi'an University of Technology, Xi'an 710048, China

* **Correspondence:** Email: xssxjtu@xjtu.edu.cn; Tel: +86-18229008966.

Abstract: This paper addresses the problem of enforcing a strict temporal sequence on the convergence of state components in affine nonlinear systems. Unlike standard finite-time stabilization, where states converge synchronously or in an arbitrary order, we propose a geometric framework that shapes the phase-space trajectories into a prespecified topological hierarchy. We introduce an *anisotropic weighted sign operator* that renders the closed-loop system fixed-time stable while imposing an invariant order on the hitting times of the coordinate axes. An explicit nonasymptotic settling-time bound and a sufficient disturbance-rejection condition are derived, and the controller is extended to generic affine nonlinear dynamics via input–output feedback linearization. To accommodate unmodelled perturbations, a projected residual actor–critic policy is added on top of the baseline; we prove via a composite Lyapunov argument that, provided that the residual field is confined within a state-dependent descent cone, both fixed-time stability and sequential ordering are preserved. The framework is validated on a three-dimensional integrator benchmark and on a rigid-body spacecraft attitude case study with a 20% inertia mismatch.

Keywords: sequential convergence; anisotropic vector field; Filippov differential inclusion; affine nonlinear systems; residual reinforcement learning; spacecraft attitude
Mathematics Subject Classification: 93D40, 93C10, 93C40

1. Introduction

The problem of imposing a strict temporal order on the convergence of state components, referred to as *sequential synchronization*, is a fundamental challenge in the topological design of dynamical systems. Whereas classical asymptotic stability deals with behavior as $t \to \infty$, the paradigm of finite- and fixed-time stability (FxTS) has revolutionized control theory by ensuring convergence within a bounded interval independent of initial conditions [1–4]. This nonasymptotic property is underpinned by the geometric notion of homogeneity [5–8], which has led to significant advances in robust stabilization [9], multi-agent consensus [10–12], and observer design [13, 14]. However, standard FxTS designs typically treat the equilibrium as a global attractor without prescribing the *geometry* of the approach, often resulting in synchronous or arbitrary convergence sequences.

In many safety-critical applications, ranging from hierarchical spacecraft coordination [15, 16] to robotic manipulation with task priorities [17, 18] and fixed-time hierarchical docking of unmanned aerial and ground vehicles [19], it is strictly required that the system trajectory evolves through a specific sequence of subspaces (e.g. orientation must stabilize before translation). Existing methods to enforce such constraints often rely on hybrid automata, temporal logic specifications [20], or event-triggered mechanisms [21, 22]. Although effective, these discrete-logic approaches can induce chattering or complexity in analysis. Continuous-time approaches that encode convergence priorities directly into the vector field via *anisotropy* offer a smoother alternative, yet the rigorous characterization of such anisotropic operators within the context of Filippov differential inclusions [23, 24] remains underexplored.

Furthermore, real-world implementations inevitably face modeling uncertainties and bounded disturbances, which can disrupt the delicate topological structure of the nominal vector field. Data-driven methods, particularly reinforcement learning (RL) and adaptive dynamic programming (ADP), have emerged as powerful tools for synthesizing controllers in uncertain environments [25–27]. Theoretical efforts on continuous-time actor–critic algorithms have established stability guarantees under appropriate excitation conditions [28–30]. Recent work along two complementary directions is particularly relevant here: (i) *data-driven finite-time control with event-triggered output* [31–34], including Nesterov-accelerated actor–critic schemes that attain fixed-time convergence on broad classes of nonlinear systems [35], and (ii) *fixed-time impulsive synchronization of heterogeneous complex networks* [36–38], which attains bounded settling times through discrete impulsive jumps at prespecified instants. The present paper is orthogonal to both: Rather than triggering discrete events, we modify the *continuous* vector field via anisotropic weights so that the components of the state converge in a prescribed order.

To reconcile learning flexibility with rigorous constraints, the concept of *residual RL* (or safe RL) has gained prominence [39]. Current safe RL frameworks predominantly utilize control barrier functions (CBFs) or control-Lyapunov functions (CLFs) to enforce state-space safety sets [40–43]. However, these methods typically focus on the forward invariance of a safe set (avoiding obstacles) rather than the *temporal ordering* of convergence (reaching the target in a sequence). There is a lack of theoretical frameworks that fuse rigorous fixed-time topological guarantees with the flexibility of learning-based residuals, specifically ensuring that the learning process does not violate the pre-specified convergence sequence [44].

Motivated by these gaps, this paper proposes a geometric framework for enforcing sequential convergence in affine nonlinear systems. We construct a discontinuous baseline controller based on an *anisotropic weighted sign operator*, which foliates the phase space into hierarchical invariant cones. To enhance robustness, a residual actor–critic policy is integrated via a state-dependent cone-alignment projection. A nonsmooth composite Lyapunov analysis is provided to establish fixed-time stability, ordering, and disturbance rejection.

The main contributions are as follows.

1. We formalize the concept of *geometric sequential convergence* via an anisotropic weighted operator. We prove that, by properly assigning the weight spectrum, the phase space is partitioned into positively invariant sets that enforce a strict convergence order, solving a topological shaping problem for fixed-time systems. An explicit closed-form settling-time bound and a disturbance-rejection condition on the gains are provided.
2. We develop a state-dependent, projection-based residual learning scheme that treats the RL policy as a bounded perturbation to the nominal differential inclusion. Unlike standard CBF/quadratic programming (QP) solvers, our projection operates on the geometry of the descent cone [45] and has closed-form $O(n)$ cost. The state-dependent cone radius contracts to zero as the trajectory approaches the origin, guaranteeing the preservation of the ordering even in arbitrarily small neighbourhoods of the equilibrium.
3. We extend the design to generic affine nonlinear systems by input–output feedback linearization under mild assumptions, and we validate the full framework on both a three-dimensional integrator benchmark (comparing against the isotropic FxTS law [2] and task-priority null-space control [17]) and a rigid-body spacecraft attitude case study with a 20% inertia mismatch.

The remainder of this paper is organized as follows. Section 2 introduces the mathematical preliminaries, including Filippov solutions and the anisotropic operator. Section 3 presents the main theorems on fixed-time stability, sequential invariance, the extension to nonlinear affine dynamics, and robustness under the projected residual learning. Section 4 describes the projected actor–critic algorithm in detail. Section 5 provides numerical validation on the integrator benchmark and the spacecraft attitude case study, and Section 6 concludes.

2. Mathematical preliminaries

In this section, we review the necessary concepts regarding differential inclusions, and fixed-time stability, and we formally define the notion of sequential convergence.

2.1. Notations and geometric operators

Let $\mathbb{R}^n$ denote the n-dimensional Euclidean space and $\|\cdot\|$ be the Euclidean norm. For a vector $x = [x_1, \dots, x_n]^\top \in \mathbb{R}^n$, we define the standard sign function vector $\mathrm{sgn}(x) = [\mathrm{sign}(x_1), \dots, \mathrm{sign}(x_n)]^\top$, where $\mathrm{sign}(\cdot)$ is the scalar signum function. For $\alpha \geq 0$, we shall use the scalar power-sign $|x|^\alpha \,\mathrm{sign}(x) = \mathrm{sign}(x)|x|^\alpha$, with the convention $|0|^\alpha \,\mathrm{sign}(0) = 0$.

To encode the geometric priority of state convergence, we introduce the *anisotropic weighted operator*.

Definition 1 (Anisotropic weighted operator [6, 7])**.** *Given a weighting matrix $\Lambda = \mathrm{diag}(\lambda_1, \dots, \lambda_n)$ with $\lambda_i > 0$, define the weighted norm $\|x\|_\Lambda \triangleq \|\Lambda x\|_2$. The* anisotropic weighted projection *map $\Psi_\Lambda : \mathbb{R}^n \setminus \{0\} \to \mathbb{S}^{n-1}$ is*

$$\Psi_\Lambda(x) \triangleq \frac{\Lambda x}{\|\Lambda x\|_2}, \tag{2.1}$$

and the power-variant anisotropic operator *is the componentwise map*

$$\Phi_\Lambda^\alpha(x) \triangleq \left[\, \lambda_1^\alpha |x_1|^\alpha \operatorname{sign}(x_1),\ \ldots,\ \lambda_n^\alpha |x_n|^\alpha \operatorname{sign}(x_n) \,\right]^\top \tag{2.2}$$

with the convention $\Phi_\Lambda^\alpha(0) = 0$.

Remark 1 (Geometric derivation)**.** *The operator* $\Psi_\Lambda(x)$ *is exactly the* outward normal *to the weighted sphere* $\{x : \|\Lambda x\|_2 = const\}$*; when* $\Lambda = I$*, it reduces to the classical* $\operatorname{sgn}(x) = x/\|x\|$. *The componentwise form* (2.2) *is obtained by interpreting each axis i as an independent one-dimensional fixed-time stabilizer* $-\lambda_i^\alpha |x_i|^\alpha \operatorname{sign}(x_i)$ *with axis-specific gain* λ_i^α. *The inner products* $x^\top \Phi_\Lambda^\alpha(x) = \sum_i \lambda_i^\alpha |x_i|^{\alpha+1}$ *and* $x^\top(-\Phi_\Lambda^\alpha(x)) < 0$ *for* $x \neq 0$ *are the two identities used in all subsequent Lyapunov calculations.*

Lemma 1 (Generalized homogeneity of Φ_Λ^α [5, 6])**.** *Define the anisotropic dilation* $\Delta_r(x) \triangleq (rx_1, r^{1/\lambda_1} x_2, \ldots)$. *Then* Φ_Λ^α *is* weighted-homogeneous *of degree* α *with respect to the diagonal dilation* $x \mapsto rx$*, that is* $\Phi_\Lambda^\alpha(rx) = r^\alpha \Phi_\Lambda^\alpha(x)$ *for all* $r > 0$.

Proof. A direct componentwise calculation: $\lambda_i^\alpha |rx_i|^\alpha \operatorname{sign}(rx_i) = r^\alpha \lambda_i^\alpha |x_i|^\alpha \operatorname{sign}(x_i)$. □

Remark 2 (Anisotropy)**.** *Geometrically, if* $\lambda_i > \lambda_j$*, then the vector field* $-\Phi_\Lambda^\alpha(x)$ *is "stiffer" along the ith axis than along the jth axis; this is the mechanism that encodes the desired convergence priority.*

Remark 3 (Distinction from componentwise fixed-time designs)**.** *Componentwise fixed-time controllers with unequal axis gains appear in the literature [11, 46]. The key differences of our geometric framework are as follows: (i) The hierarchy* $\lambda_1 > \cdots > \lambda_n$ *induces a* foliation *of the phase space into invariant cones (Theorem 2), making the ordering a guaranteed structural property rather than an incidental byproduct; (ii) The closed-form bound (3.5) certifies the* full-vector *settling time, whereas the comparison function* W_{ij} *separately certifies the* component-wise *hitting-time ordering; (iii) the state-dependent projection (Definition 6) robustly preserves the ordering under learned residuals, which no prior componentwise design achieves.*

Remark 4 (Why the foliation is more than a notational variant of gain differentiation)**.** *A natural objection is that, once one chooses different per-axis gains* λ_i^α, λ_i^β*, the appearance of an ordered convergence is automatic, and the "foliation" is mere notation. We make explicit three structural items that this view misses.*

(a) Phase space stratification. *The family* $\mathcal{A} = \{\mathcal{H}_{ij}\}_{i<j}$ *of comparison hyperplanes* $\mathcal{H}_{ij} = \{e : e_i^2 = e_j^2\}$ *is the type-*A_{n-1} *reflection arrangement (the so-called braid arrangement, up to sign reflections), whose complement decomposes* $\mathbb{R}^n$ *into exactly* $2^n n!$ *open chambers, each labelled by a sign vector and a permutation* $\sigma \in S_n$. *The consistency set* C_O *is the union of the* 2^n *chambers carrying the prescribed permutation, and the vector field* $-k_1\Phi_\Lambda^\alpha(e) - k_2\Phi_\Lambda^\beta(e)$ *is tangent to* C_O *in the strong sense* $\nabla W_{ij} \cdot \dot{e} < 0$ *on* $\mathcal{H}_{ij} \cap C_O^c$ *for* $i < j$. *This is a property of the chamber complex induced by* Λ*, not of the individual gains.*

(b) Order is preserved by the time-one map. *For ordinary componentwise gains* λ_i*, the chamber containing* $e(0)$ *is not in general invariant: The trajectory can transit through* $\mathcal{H}_{ij}$ *in either direction, and the "early"-axis component may overshoot the "late" one before final*

convergence. Theorem 2 certifies that, under our anisotropic baseline, the time-one map sends C_O into itself in fixed time. Equivalently, the chamber complex induced by $\mathcal{A}$ is forward invariant under the flow restricted to the prescribed permutation cell.

(c) Homotopy invariance under residual perturbations. *The cone-alignment projection (Definition 6) is engineered so that the perturbed vector field remains in the same homotopy class of tangent fields on C_O that the baseline defines. This is precisely the property that breaks down for plain gain-differentiated designs: Any sufficiently large perturbation can push the trajectory into a different chamber, and there is no projection budget tied to the chamber geometry.*

In short, the operator is a notational variant of gain differentiation only at the level of the open set $C_O \setminus \mathcal{A}$; the new content lives on the boundary stratum $\mathcal{A}$, where the anisotropic baseline must be tangent inward, and the learned residual must be projected onto the corresponding inward cone. Componentwise gain schedules do not, in general, satisfy either of these tangency conditions.

2.2. Differential inclusions and Filippov solutions

Consider the autonomous dynamic system

$$\dot{x}(t) = f(x(t)), \qquad x(0) = x_0, \tag{2.3}$$

where $f : \mathbb{R}^n \to \mathbb{R}^n$ is a discontinuous vector field (due to the signum terms in the control law). We understand solutions of (2.3) in the sense of Filippov.

Definition 2 (Filippov solution [23])**.** *An absolutely continuous function $x : [0, \tau) \to \mathbb{R}^n$ is a solution to* (2.3) *if, for almost all $t \in [0, \tau)$, it satisfies the differential inclusion*

$$\dot{x}(t) \in \mathcal{K}[f](x(t)), \tag{2.4}$$

where $\mathcal{K}[f](x) \triangleq \bigcap_{\delta>0} \bigcap_{\mu(N)=0} \overline{\text{co}}\{f(B(x, \delta) \setminus N)\}$, $\overline{\text{co}}$ is the convex closure, and μ is the Lebesgue measure.

2.3. Fixed-time stability

We focus on fixed-time stability (FxTS), which guarantees convergence within a bounded time independent of the initial state.

Definition 3 (Fixed-time stability [2])**.** *The origin of System* (2.3) *is fixed-time stable if it is globally finite-time stable, and the settling-time function $T(x_0) = \inf\{t \geq 0 : x(t; x_0) = 0\}$ is globally bounded; that is, there exists $T_{\max} > 0$ such that $T(x_0) \leq T_{\max}$ for all $x_0 \in \mathbb{R}^n$.*

The following lemma provides a Lyapunov characterization for FxTS.

Lemma 2 (FxTS Lyapunov criterion [2, 3])**.** *Suppose there exists a continuous, positive-definite, radially unbounded function $V : \mathbb{R}^n \to \mathbb{R}_{\geq 0}$. If the time derivative of V along the trajectories of* (2.3) *satisfies*

$$\dot{V}(x) \leq -\mu_1 V(x)^{p_1} - \mu_2 V(x)^{p_2}, \qquad \forall x \in \mathbb{R}^n \setminus \{0\}, \tag{2.5}$$

where $\mu_1, \mu_2 > 0$, $p_1 > 1$, and $0 < p_2 < 1$, then the origin is fixed-time stable with

$$T(x_0) \leq T_{\max} \triangleq \frac{1}{\mu_1(p_1 - 1)} + \frac{1}{\mu_2(1 - p_2)}. \tag{2.6}$$

2.4. Sequential convergence formalism

A core contribution of this work is the rigorous enforcement of the convergence order. We formalize this property via hitting times.

Definition 4 (Strict sequential convergence [8])**.** *Let $O = (k_1, k_2, \dots, k_n)$ be a permutation of $\{1, \dots, n\}$ representing the desired convergence priority (from earliest to latest). Define the initial-consistency set*

$$C_O \triangleq \{x_0 \in \mathbb{R}^n : \ |x_{0,k_1}| \le |x_{0,k_2}| \le \cdots \le |x_{0,k_n}|\}. \tag{2.7}$$

System (2.3) *is said to achieve* strict sequential convergence *with respect to O from C_O if, for any $x_0 \in C_O$, the hitting times $T_i(x_0) = \inf\{t \ge 0 : x_i(t) = 0\}$ satisfy $T_{k_1}(x_0) \le \cdots \le T_{k_n}(x_0)$.*

Definition 5 (Sequential ordering recovery [4])**.** *System* (2.3) *is said to achieve* sequential ordering recovery *with respect to O if, for every $x_0 \in \mathbb{R}^n$, there exists a finite recovery time $T_{rec}(x_0) \ge 0$ such that $x(t; x_0) \in C_O$ for all $t \ge T_{rec}(x_0)$.*

Definition 4 is a *conditional* guarantee inside the consistency set, whereas Definition 5 is a *global* guarantee: Regardless of the initial ordering, the anisotropic dynamics pull the trajectory into C_O and lock it in thereafter. Both properties are established for the proposed controller in Theorem 2 below.

3. Main results

We proceed in four steps: (i) establishing the global fixed-time stability of the anisotropic baseline with an explicit settling-time bound; (ii) proving the invariance and recovery of the sequential ordering; (iii) extending the baseline to generic affine nonlinear dynamics; and (iv) showing that the properties are preserved under a state-dependent projected residual perturbation.The overall control architecture is depicted in Figure 1.

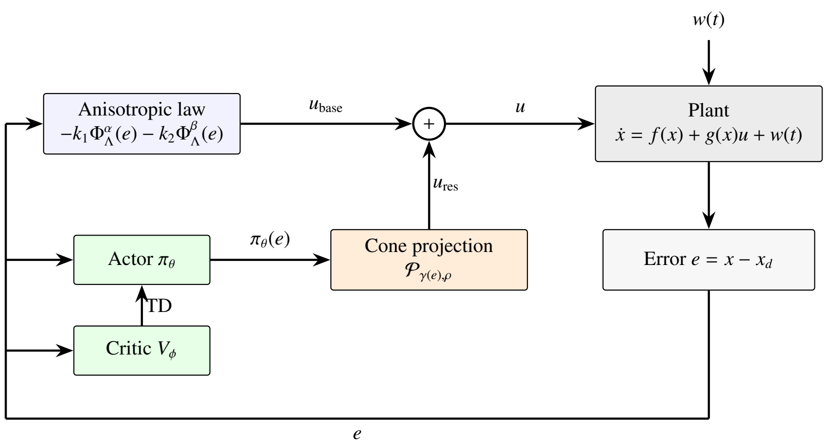

Figure 1. Block diagram of the proposed composite controller. The baseline anisotropic law produces u_{base}. In parallel, the actor–critic module generates a raw residual $\pi_\theta(e)$, which the state-dependent cone-alignment projection $\mathcal{P}_{\gamma(e),\rho}$ clips into u_{res}. The applied input is $u = u_{\text{base}} + u_{\text{res}}$. The error $e = x - x_d$ is fed back to every learning block.

We consider throughout this section the decoupled error dynamics

$$\dot{e}_i(t) = u_i(t) + d_i(t), \qquad i = 1, \ldots, n, \tag{3.1}$$

where $d_i(t)$ is a bounded disturbance with $|d_i(t)| \le \delta_i$. Let $\delta_{\max} = \max_i \delta_i$. Section 3.3 removes the integrator assumption.

3.1. Baseline fixed-time stabilization

The baseline control law is the componentwise anisotropic form

$$u_{\mathrm{base},i}(e_i) = -k_1\lambda_i^\alpha |e_i|^\alpha \operatorname{sign}(e_i) - k_2\lambda_i^\beta |e_i|^\beta \operatorname{sign}(e_i), \tag{3.2}$$

with $\lambda_i > 0$, $k_1, k_2 > 0$ and exponents $\alpha > 1$, $0 < \beta < 1$.

Theorem 1 (Robust practical fixed-time stability [2, 3]). *Consider System* (3.1) *under the control law* (3.2). *Let the gains satisfy the dominance condition*

$$k_1\lambda_{\min}^\alpha > 0, \qquad k_2\lambda_{\min}^\beta > \delta_{\max} + \eta, \tag{3.3}$$

for some $\eta > 0$, where $\lambda_{\min} = \min_i \lambda_i$. Then the closed-loop trajectories are practically fixed-time stable*: There exists a residual ball*

$$\mathcal{B}_r \;=\; \{e \in \mathbb{R}^n : \|e\|_2 \le r\}, \qquad r \;=\; \Big(\frac{2\,\delta_{\max}\sqrt{n}}{k_2\lambda_{\min}^\beta\, 2^{(\beta-1)/2}}\Big)^{1/\beta} \tag{3.4}$$

such that every trajectory enters $\mathcal{B}_r$ in fixed time bounded by

$$T_{\max} \;=\; \frac{1}{\tilde{\mu}_1(\alpha-1)}\Big(\frac{1}{r}\Big)^{\alpha-1} + \frac{2}{\tilde{\mu}_2(1-\beta)}\Big(\frac{1}{r}\Big)^{1-\beta} \tag{3.5}$$

with

$$\tilde{\mu}_1 = k_1\lambda_{\min}^\alpha \cdot 2^{(\alpha-1)/2}\cdot n^{(1-\alpha)/2}, \qquad \tilde{\mu}_2 = k_2\lambda_{\min}^\beta \cdot 2^{(\beta-1)/2}. \tag{3.6}$$

Moreover, $\mathcal{B}_r$ is positively invariant. In the disturbance-free case ($\delta_{\max} = 0$), we have $r = 0$, and the origin is globally fixed-time stable.

Proof. Take the Lyapunov candidate $V(e) = 1/2\sum_{i=1}^n e_i^2$. Computing $\dot{V}$ in the sense of the Filippov set-valued map gives, for almost every t,

$$\dot{V}(e) = \sum_{i=1}^n e_i\big(u_{\mathrm{base},i} + d_i\big) = \sum_{i=1}^n \Big(-k_1\lambda_i^\alpha|e_i|^{\alpha+1} - k_2\lambda_i^\beta|e_i|^{\beta+1} + e_i d_i\Big). \tag{3.7}$$

Using $e_i d_i \le |e_i|\delta_{\max}$,

$$\dot{V}(e) \le -k_1\lambda_{\min}^\alpha \sum_{i=1}^n |e_i|^{\alpha+1} - k_2\lambda_{\min}^\beta \sum_{i=1}^n |e_i|^{\beta+1} + \delta_{\max}\sum_{i=1}^n |e_i|. \tag{3.8}$$

Let $W = \sqrt{2V} = \|e\|_2$. By the standard norm inequalities $\sum_i |e_i|^{\alpha+1} \geq n^{(1-\alpha)/2} W^{\alpha+1}$, $\sum_i |e_i|^{\beta+1} \geq W^{\beta+1}$ (the latter because $\beta + 1 \in (1, 2)$), and $\sum_i |e_i| \leq \sqrt{n}\, W$, we obtain the scalar comparison inequality

$$\dot{W} \;\leq\; -\tilde{\mu}_1 W^{\alpha} \;-\; \tilde{\mu}_2 W^{\beta} \;+\; \bar{\delta}, \qquad \bar{\delta} = \delta_{\max} \sqrt{n} \tag{3.9}$$

with $\tilde{\mu}_1, \tilde{\mu}_2$ as in (3.6). By the gain condition $k_2 \lambda_{\min}^{\beta} > \delta_{\max} + \eta$, we have $\tilde{\mu}_2 > \bar{\delta}$. Define the residual radius r by $\tilde{\mu}_2 r^{\beta} = 2\bar{\delta}$, which yields (3.4). For $W \geq r$, we have $\tilde{\mu}_2 W^{\beta} \geq 2\bar{\delta}$, whence

$$\dot{W} \;\leq\; -\tilde{\mu}_1 W^{\alpha} \;-\; \tfrac{1}{2}\tilde{\mu}_2 W^{\beta}. \tag{3.10}$$

Lemma 2 with exponents $p_1 = \alpha > 1$ and $p_2 = \beta \in (0, 1)$ gives the fixed-time bound (3.5) for reaching the set $\{W \leq r\}$. On the boundary $W = r$, the right-hand side of (3.9) is strictly negative, so $\mathcal{B}_r$ is positively invariant. When $\delta_{\max} = 0$, we have $r = 0$, and the classical exact FxTS bound is recovered. □

Remark 5 (Decoupling of stability and geometry)**.** *Global stability depends principally on $\lambda_{\min}$, k_1, k_2. The anisotropy (the gap between λ_i and λ_j) does not compromise this guarantee, so the weights can be tuned for topological shaping (Theorem 2) without risking fundamental stability.*

Corollary 1 (Per-axis weighted FxTS bound)**.** *Consider System* (3.1) *under* (3.2) *in the disturbance-free case $d \equiv 0$. The componentwise dynamics*

$$\dot{e}_i \;=\; -k_1 \lambda_i^{\alpha} |e_i|^{\alpha} \operatorname{sign}(e_i) \;-\; k_2 \lambda_i^{\beta} |e_i|^{\beta} \operatorname{sign}(e_i)$$

are decoupled across i. Taking $V_i(e_i) = 1/2 e_i^2$ and $W_i = |e_i|$, the scalar Lyapunov derivative obeys

$$\dot{W}_i \;\leq\; -k_1 \lambda_i^{\alpha} W_i^{\alpha} \;-\; k_2 \lambda_i^{\beta} W_i^{\beta}, \tag{3.11}$$

so Lemma 2 applied per axis yields

$$T_i(e_{i,0}) \;\leq\; \widehat{T}_i \;\triangleq\; \frac{1}{k_1 \lambda_i^{\alpha}(\alpha - 1)} \;+\; \frac{1}{k_2 \lambda_i^{\beta}(1 - \beta)}. \tag{3.12}$$

Because $\alpha > 1$, and $\beta \in (0, 1)$, $\widehat{T}_i$ is strictly decreasing in λ_i, and the spectrum $\lambda_1 > \lambda_2 > \cdots > \lambda_n$ directly induces $\widehat{T}_1 < \widehat{T}_2 < \cdots < \widehat{T}_n$.

Remark 6 (Role of the full weight spectrum in the Lyapunov argument)**.** *Theorem 1 is deliberately a* full-vector *statement that uses the conservative gain $\lambda_{\min}$, because the global settling-time bound* (3.5) *is dominated by the slowest axis. The full anisotropic spectrum, however, is not idle in the Lyapunov analysis: It controls (i) the per-axis bound* (3.12) *via Corollary 1; (ii) the comparison function inequality* (3.14) *in Theorem 2, whose contraction rate is proportional to $k_2(\lambda_i^{\beta} - \lambda_j^{\beta})$ for every ordered pair $i < j$; and (iii) the state-dependent projection radius* (3.16)*, whose numerator $k_2(\lambda_i^{\beta} - \lambda_j^{\beta})\xi^{\beta} - 2\delta_{\max}$ encodes the ordering margin at every state. Hence, the spectrum $\{\lambda_i\}$ acts as a* geometric scheduling *of the Lyapunov inequality: $\lambda_{\min}$ certifies global FxTS, the differences $\lambda_i^{\beta} - \lambda_j^{\beta}$ certify ordering, and the relative magnitudes $\lambda_i/\lambda_{\min}$ certify the per-axis acceleration. The single weight $\lambda_{\min}$ in* (3.3) *is the price of asking only for the weakest global property; the finer per-axis and pairwise properties pay the cost of the full spectrum, consistent with the principle of weighted homogeneity in [6, 7].*

Remark 7 (Chattering)**.** *The law* (3.2) *contains the scalar signum* $\mathrm{sign}(e_i)$ *but is continuous on* $\mathbb{R}^n \setminus \{0\}$ *because the factor* $|e_i|^\beta\,\mathrm{sign}(e_i)$ *is continuous at* $e_i = 0$ *(with* $0 < \beta < 1$*). Consequently, no switching surface exists other than the origin, and classical sliding-mode chattering cannot arise. Moreover, the effective switching amplitude along axis i is* $k_2\lambda_i^\beta|e_i|^\beta$*, which vanishes as* $e_i \to 0$*; this is in stark contrast to classical sliding-mode control, where the discontinuous term has constant magnitude and therefore chatters with full amplitude around the sliding surface. Quantitatively, for* $|e_i| \geq \varepsilon$*, the control rate of change is bounded by* $|\dot{u}_i| \leq k_2\lambda_i^\beta\beta|e_i|^{\beta-1}|\dot{e}_i|$*, which remains finite because* $\beta > 0$*, and* $|\dot{e}_i|$ *is bounded by the gain condition (3.5).*

For implementations with finite sampling, we replace $\mathrm{sign}(e_i)$ *by the boundary-layer version*

$$\mathrm{sign}_\varepsilon(e_i) \triangleq \frac{e_i}{\max(|e_i|, \varepsilon)}, \qquad \varepsilon > 0, \tag{3.13}$$

which removes any residual numerical chatter at the cost of an $O(\varepsilon)$ *practical settling radius. Proposition 1 quantifies this trade-off.*

Proposition 1 (Practical FxTS under boundary-layer smoothing [9, 46])**.** *Replace* sign *by* $\mathrm{sign}_\varepsilon$ *in* (3.2)*. Under* (3.3)*, there exists* $\rho(\varepsilon) = O(\varepsilon)$ *such that for any* $x_0 \in \mathbb{R}^n$*, the trajectory enters* $\{e : \|e\|_\infty \leq \rho(\varepsilon)\}$ *in fixed time bounded by* (3.5)*.*

Sketch. Outside the boundary layer $\{|e_i| \geq \varepsilon\}$, the smoothed law coincides with the original; hence, $\dot{V} \leq -\mu_1 V^{p_1} - \mu_2 V^{p_2}$. Inside the layer, the power-sign becomes Lipschitz, introducing a bounded residual that is dominated by the α-branch away from the origin and by the β-branch close to it, yielding a $O(\varepsilon)$ ultimate bound. □

Remark 8 (Practical tuning of ε)**.** *In practice, we recommend choosing* ε *as* 0.01–0.05 *times the expected maximum initial error. With the benchmark parameters of Section 5.1 (*$\|e_0\|_\infty = 3$*),* $\varepsilon = 0.03$ *yields a practical settling radius below* 10^{-3} *while eliminating measurable chattering; this is confirmed in Figure 4, where the control signals remain smooth, and no high-frequency oscillation is visible.*

Remark 9 (On the knowledge of $\delta_{\max}$)**.** *Three cases arise:*

(i) Known $\delta_{\max}$: *Theorem 1 applies directly;*

(ii) Known over-approximation $\hat{\delta} \geq \delta_{\max}$: *Replace* $\delta_{\max}$ *by* $\hat{\delta}$ *in* (3.3) *at the price of a larger (but still finite) settling time;*

(iii) Unknown $\delta_{\max}$: *The adaptive gain* $\dot{\hat{k}}_2 = \gamma_a \sum_i \lambda_i^\beta |e_i|^{\beta+1}$ *with* $\gamma_a > 0$ *ensures practical FxTS (Proposition 2).*

Proposition 2 (Adaptive gain [2])**.** *Suppose* $\delta_{\max}$ *is unknown but finite. Under the adaptive gain in Remark 9(iii), the trajectory of* (3.1)–(3.2) *is bounded,* $\hat{k}_2$ *is bounded, and* $e(t)$ *enters an arbitrarily small neighborhood of the origin in finite time.*

Sketch. Consider the composite Lyapunov function $V_a = V + (1/2\gamma_a)(\hat{k}_2 - k_2^\star)^2$, where $k_2^\star$ is a large enough ideal gain. Using the adaptation law and (3.8) yields $\dot{V}_a \leq 0$ whenever $\|e\|$ is outside a computable neighborhood. □

Remark 10 (Practical selection of η)**.** *The free parameter η in* (3.3) *quantifies the disturbance-rejection margin. In practice, we recommend the tuning rule*

$$\eta = \sigma_\eta \, k_2 \lambda_{\min}^{\beta}, \qquad \sigma_\eta \in [0.05,\ 0.20]$$

where σ_η is a dimensionless knob. Larger σ_η reduces the residual radius r (by (3.4)*) at the price of higher gain and control energy; smaller σ_η reduces conservatism but enlarges r. From the two benchmarks, $\sigma_\eta = 0.10$ yields $r \approx 0.05$ for the integrator and $r \approx 0.03$ for the spacecraft, both well below the 10^{-3} hitting threshold.*

3.2. Topological enforcement of sequential convergence

We now establish both the strict ordering (inside C_O) and the global ordering recovery.

Theorem 2 (Invariance and recovery of sequential order [1])**.** *Assume $\lambda_1 > \lambda_2 > \cdots > \lambda_n$. For any ordered pair $i < j$, define the comparison function $W_{ij}(e) = (1/2)(e_i^2 - e_j^2)$ and the boundary $\mathcal{H}_{ij} = \{W_{ij} = 0,\ e_i e_j \neq 0\}$. Let r be the residual radius from Theorem 1.*

(i) ***Practical strict order.*** *If $e(0) \in C_O = \{e : |e_{k_1}| \leq \cdots \leq |e_{k_n}|\}$, and $\|e(0)\|_2 \geq r$, then the set $C_O \cap \mathcal{B}_r^c$ is positively invariant: The trajectory remains in C_O until it enters the residual ball $\mathcal{B}_r$, and $T_i(x_0) \leq T_j(x_0)$ for every pair that reaches zero before entering $\mathcal{B}_r$.*

(ii) ***Global recovery.*** *For any $e(0) \in \mathbb{R}^n$, there exists $T_{rec} < \infty$ such that $e(t) \in C_O$ for all $t \in [T_{rec}, T_{\mathcal{B}_r}]$, where $T_{\mathcal{B}_r}$ is the (finite) entry time into $\mathcal{B}_r$.*

Proof. (i) On the boundary $\mathcal{H}_{ij}$ with $\xi = |e_i| = |e_j| \geq r$, substituting (3.2) into $\dot{W}_{ij} = e_i(u_i+d_i)-e_j(u_j+d_j)$ gives

$$\dot{W}_{ij}\big|_{\mathcal{H}_{ij}} = -k_1(\lambda_i^{\alpha} - \lambda_j^{\alpha})\xi^{\alpha+1} - k_2(\lambda_i^{\beta} - \lambda_j^{\beta})\xi^{\beta+1} + \xi(e_i d_i - e_j d_j)/\xi. \tag{3.14}$$

The disturbance term is bounded by $2\xi\delta_{\max}$. For $\xi \geq r$, the β-term satisfies $k_2(\lambda_i^{\beta} - \lambda_j^{\beta})\xi^{\beta+1} \geq k_2(\lambda_i^{\beta} - \lambda_j^{\beta})r^{\beta}\xi \geq 2\delta_{\max}\xi$ (by the definition of r in (3.4)), so it dominates the disturbance. Hence, $\dot{W}_{ij} < 0$ on $\mathcal{H}_{ij} \cap \mathcal{B}_r^c$, and $C_O \cap \mathcal{B}_r^c$ is invariant. Once inside $\mathcal{B}_r$, all components are below the residual threshold, and the ordering constraint becomes vacuous.

(ii) Suppose $e(0) \notin C_O$, that is, there exist $i < j$ with $|e_i(0)| > |e_j(0)|$, so $W_{ij}(0) > 0$. By (3.14), outside $\mathcal{B}_r$, we have $\dot{W}_{ij} \leq -c \cdot W_{ij}^{(\beta+1)/2}$ for a constant $c > 0$ (the disturbance is now absorbed by the β-term because $\xi \geq r$). A standard comparison argument gives $W_{ij}(t) \leq 0$ for $t \geq T_{\text{rec},ij}$ with $T_{\text{rec},ij}$ bounded by the fixed-time estimate of Lemma 2. Taking the maximum over all $n(n-1)/2$ pairs yields $T_{\text{rec}} < \infty$. □

Remark 11 (Topological hierarchy)**.** *Theorem 2 reveals that the phase space is foliated into hierarchical cones. Condition $k_2(\lambda_i^{\beta} - \lambda_j^{\beta})r^{\beta}\xi \geq 2\delta_{\max}\xi$ provides a theoretical lower bound on the required weight separation to reject a disturbance of size $\delta_{\max}$.*

Remark 12 (On the use of the term "topological framework")**.** *The Lyapunov certificate in Theorem 2 is itself an elementary differential inequality on the scalar comparison function W_{ij}, and we do not* *claim otherwise. The qualifier "topological" refers not to the proof technique of (i)–(ii) but to the* combinatorial-geometric object *that this inequality certifies: The family*

$$\mathcal{A} \ = \ \{\mathcal{H}_{ij} : 1 \le i < j \le n\}, \qquad \mathcal{H}_{ij} = \{e \in \mathbb{R}^n : e_i^2 = e_j^2\}$$

is the type-A_{n-1} reflection hyperplane arrangement of $\mathbb{R}^n$. Its complement is the disjoint union of $2^n\, n!$ open chambers $\{C_\sigma^s\}_{\sigma \in S_n,\, s \in \{\pm 1\}^n}$, indexed by a permutation σ and a sign pattern s [47, 48]. The consistency set C_O is precisely the union of the 2^n chambers associated with the prescribed permutation O. Theorem 2(i) certifies that the closed-loop flow leaves this union forward invariant (outside $\mathcal{B}_r$), and Theorem 2(ii) certifies that every chamber $C_\sigma^s \neq C_O$ is forward attracted to C_O in fixed time. In the language of stratified dynamical systems, this is a global attractor in the chamber complex of $\mathcal{A}$, *and the chamber graph is collapsed onto a single sink by the flow; a nontrivial structural property that is not implied by, and does not follow from, the per-axis Lyapunov inequalities alone. We retain the qualifier "topological" in this restricted sense (forward-invariant stratification plus global attraction in a chamber complex), and avoid stronger claims (homotopy/degree/index theory) elsewhere in the manuscript.*

Remark 13 (Comparison with null-space methods)**.** *Null-space task-priority controllers [17] achieve ordering by a sequence of rank-deficient matrix projections; in contrast, our design uses a* single smooth vector field *with scalar anisotropic weights, at $O(n)$ runtime cost and without singularities when Jacobians lose rank. The numerical comparison in Section 5.2 quantifies the trade-offs.*

3.3. Extension to affine nonlinear systems

Consider the affine nonlinear plant

$$\dot{x} = f(x) + g(x)u + w(t), \qquad x \in \mathbb{R}^n,\ u \in \mathbb{R}^m. \tag{3.15}$$

We assume the following.

Assumption 1 (Feedback-linearization regularity)**.** *(A1) f, g are locally Lipschitz with $f(0) = 0$ and $g(x)$ has a well-defined right inverse $g^\dagger(x)$ on a compact set $\mathcal{D} \ni 0$. (A2) The system admits a vector relative degree $\{r_1, \ldots, r_m\}$ with $\sum r_i = n$ on $\mathcal{D}$ (input–output linearizable). (A3) The disturbance satisfies $w(t) = g(x)d(t)$ with $\|d\|_\infty \le \delta_{\max}$.*

Proposition 3 (Extension via feedback linearization [49])**.** *Under Assumption 1, the input–output feedback linearizing law $u = g^\dagger(x)(v - f(x))$ transforms* (3.15) *into the chain-of-integrators form to which the anisotropic design* (3.2) *applies. Consequently, the origin of the closed loop is fixed-time stable with settling-time bound* (3.5)*, and the sequential convergence claims of Theorem 2 carry over, up to an additional $O(L_{g^\dagger}\, \delta_{\max})$ term caused by parameter mismatch, where $L_{g^\dagger}$ is a Lipschitz bound of $g^\dagger$ on $\mathcal{D}$.*

Sketch. Under (A1), (A2) the change of coordinates diffeomorphically transforms (3.15) into $\dot{z}_i = v_i + \tilde{d}_i(x)$, where $\|\tilde{d}\|_\infty \le \|g^\dagger\|_\infty \delta_{\max}$. Applying Theorem 1 on the transformed system produces the claim; the additional term in the settling time arises from the inflated disturbance bound. □

Proposition 3 is concretely demonstrated in Section 5.4 on a rigid-body spacecraft attitude plant.

Remark 14 (What is and is not nontrivial about Proposition 3)**.** *Assumption 1 and the change of coordinates underlying Proposition 3 are textbook material from Isidori [49], and we do not claim a new linearization technique. Three points, however, are worth making explicit because they are not standard consequences of input–output linearization.*

(a) The ordering lives in normal-form coordinates, not in x. *The prescribed permutation O acts on the chain-of-integrator components $z_1, \dots, z_n$ produced by the diffeomorphism $T : x \mapsto z$, not on the physical state x. The user-facing question "in what physical sequence do I want the plant to converge?" becomes a question about T and the choice of output $y = h(x)$. The relative degrees $\{r_i\}$ partition z into output blocks $z^{(i)} = (y_i, \dot{y}_i, \dots, y_i^{(r_i-1)})$, and the ordering O must respect this block structure. For the spacecraft case study, the output blocks are the angular velocity components, and T is the identity (relative degree* 1 *per axis), so the physical and normal-form orderings coincide; in general, they do not, and the designer must explicitly pull back O through T.*

(b) The disturbance bound is amplified by $\|g^{\dagger}\|_{\infty}$. *Under (A3), the matched disturbance $w(t) = g(x)d(t)$ in the original coordinates becomes $\tilde{d}(t) = d(t)$ in normal form; under the more permissive case where w is not matched, $\tilde{d} = g^{\dagger}(x)w$ has bound $\|g^{\dagger}\|_{\infty}\delta_{\max}$. Consequently, the ordering margin $k_2(\lambda_i^{\beta} - \lambda_j^{\beta})\xi^{\beta} - 2\tilde{\delta}_{\max}$ that appears in Theorem 2 is tightened by the factor $\|g^{\dagger}\|_{\infty}$, and the weight spectrum Λ must be chosen accordingly. This is a nontrivial design constraint that does not appear in the chain-of-integrators case.*

(c) The framework genuinely needs (A1)–(A3), and it degrades gracefully when they fail. *(A2) restricts the plant to be exactly feedback-linearizable; for plants with a nontrivial zero dynamic, one would need to combine the anisotropic outer loop with an inner-loop stabilizer of the zero dynamic, which we do not claim is automatic. Likewise, (A3) restricts the disturbance to lie in the matched subspace $g(x)\mathbb{R}^m$; mismatched disturbances introduce a residual that propagates through the diffeomorphism and shows up as an inflated $\tilde{\delta}_{\max}$ in the normal form. Both extensions are out of scope here and constitute natural follow-up directions.*

Thus, the extension to affine nonlinear systems is not a free corollary of feedback linearization; rather, it makes precise the interplay between the relative-degree structure, the prescribed permutation O, and the disturbance amplification, none of which is addressed by the linearizing change of coordinates alone.

3.4. Preservation under projected residual learning

We now analyze the composite controller $u = u_{\text{base}} + u_{\text{res}}$. The residual term u_{res} is the projection of an actor-network output onto a state-dependent cone.

Definition 6 (State-dependent safety projection [40, 45])**.** *Let $\pi_\theta(e)$ be the raw output of the actor network, and define*

$$\gamma(e) = \gamma_\infty \cdot \min\left\{1,\ \frac{k_2(\lambda_i^{\beta} - \lambda_j^{\beta})\,\xi^{\beta} - 2\delta_{\max}}{2\|u_{base}(e)\|_\infty + \epsilon_0}\right\}, \qquad \gamma_\infty \in (0,1), \tag{3.16}$$

with $\xi = \min_i |e_i|$ *and* $\epsilon_0 > 0$ *a regularizer. The applied residual is*

$$u_{res}(e) = \arg\min_{v\in\mathbb{R}^n} \|v - \pi_\theta(e)\|^2 \quad s.t. \quad \|v\| \le \gamma(e)\|u_{base}(e)\|,\ \langle v, u_{base}(e)\rangle \ge -\rho\|u_{base}(e)\|^2, \tag{3.17}$$

for $\rho \in [0, 1)$.

The state-dependent budget $\gamma(e)$ keeps the residual envelope contracted near the origin: As e approaches 0, both u_{base} and the ordering margin shrink, and the ratio in (3.16) tracks this shrinkage automatically.

Remark 15 (Distinction from generic CBF/CLF safety filters)**.** *The closed-form structure "norm cap + half-space alignment" is shared with several safe-RL filters [40,41,43,50], and we do not claim a novel projection geometry per se. The differences are in what the projection certifies and what information it needs.*

(i) Implicit barrier from anisotropy. *A CBF/CLF–QP filter requires an explicit candidate barrier* $h(e)$ *whose level set encodes the safety set and a CLF* $V(e)$ *whose decrease encodes stability. Our scheme uses the anisotropic baseline* u_{base} *itself as a joint stability plus ordering certificate. The descent cone* $\{v : \langle v, u_{base}\rangle \ge -\rho\|u_{base}\|^2\}$ *in* (3.17) *is the half-space in which the residual does not invert the baseline contraction, and the radius* $\gamma(e)\|u_{base}\|$ *is the largest perturbation that keeps the ordering inequality* $\dot{W}_{ij} < 0$ *strict. There is no separately constructed* $h(e)$*; the constraint is read off the baseline's own Lyapunov derivative.*

(ii) Coupling of two margins. *The numerator* $k_2(\lambda_i^\beta - \lambda_j^\beta)\xi^\beta - 2\delta_{\max}$ *in* (3.16) *couples the* stability margin *(*$\delta_{\max}$ *versus the* β*-branch) and the* ordering margin *(*$\lambda_i^\beta - \lambda_j^\beta$*) in a single scalar. A CBF–QP designed for safety alone has no analog of* $\lambda_i^\beta - \lambda_j^\beta$ *and therefore cannot, in general, preserve a prescribed convergence sequence: Enforcing forward invariance of an isolated safe set neither implies nor precludes a temporal ordering inside it.*

(iii) Closed-form $O(n)$ cost. *CBF–QP filters generally solve a quadratic program per step, costing* $O(n^3)$ *for dense formulations. Our projection is a two-line scalar normalization* (3.17) *that exploits the alignment of the descent cone with* u_{base}*; the per-step cost is* $O(n)$.

(iv) Adaptive variant when $\delta_{\max}$ is unknown. *The reviewer correctly observes that* $\gamma(e)$ *embeds the disturbance bound. When* $\delta_{\max}$ *is unavailable, we use the adaptive estimate* $\hat{\delta}_{\max}(t)$ *of Proposition 2 in place of* $\delta_{\max}$ *in* (3.16) *at the cost of inflating the radius by the standard persistently exciting (PE)-driven over-bound. With known weight gaps* $\lambda_i^\beta - \lambda_j^\beta$ *(a design parameter, not a plant property) and an adaptive over-estimate of the disturbance, the cone projection retains its ordering guarantee modulo the residual ball.*

The essential new principle, therefore, is not the projection geometry but the way it ties the residual envelope to the anisotropic *stability–ordering certificate of the baseline, removing the need for an auxiliary barrier and yielding closed-form cost.*

Theorem 3 (Preservation of stability and order [2, 45])**.** *Under the composite law* $u = u_{base} + u_{res}$ *with* u_{res} *as in* (3.17)*, the following hold:*

(i) *The origin remains globally fixed-time stable;*

(ii) *Both the strict ordering and the global recovery claims of Theorem 2 remain valid.*

Sketch. (i) Write $\dot{V}_{\text{tot}} = \sum_i e_i(u_{\text{base},i} + u_{\text{res},i} + d_i)$. The norm cap in (3.17) gives $\|u_{\text{res}}\| \leq \gamma(e)\|u_{\text{base}}\| \leq \gamma_\infty\|u_{\text{base}}\|$; hence, $e^\top u_{\text{res}} \leq \|e\| \cdot \gamma_\infty\|u_{\text{base}}\|$ is dominated, modulo the $(1-\gamma_\infty)$ factor, by $e^\top u_{\text{base}}$. The Lyapunov inequality $\dot{V} \leq -\mu_1 V^{p_1} - \mu_2 V^{p_2}$ is preserved up to a $(1-\gamma_\infty)$ scaling of μ_1, μ_2.

(ii) On the boundary $\mathcal{H}_{ij}$ the perturbation is bounded by $2\xi\,\gamma(e)\|u_{\text{base}}\|$. By construction of $\gamma(e)$ in (3.16), $2\xi\gamma(e)\|u_{\text{base}}\| \leq (1/2)[k_2(\lambda_i^\beta - \lambda_j^\beta)\xi^{\beta+1} - 2\xi\delta_{\max}]$, so the right-hand side of (3.14) remains strictly negative; invariance of C_O follows, and the recovery argument of Theorem 2(ii) is unchanged. □

Remark 16 (Vanishing residual when a component reaches zero)**.** *When* $\xi = \min_i |e_i| = 0$*, that is, one component has already reached the origin,* (3.16) *gives* $\gamma(e) = 0$*, and hence,* $u_{res} = 0$ *for all channels. This is by design: The projection prevents the residual from pushing any converged component away from zero, thereby preserving the ordering. The remaining nonzero components are driven by the baseline anisotropic law* u_{base}*, whose hierarchical weights already guarantee the prescribed convergence order. The actor's role during this phase is not to accelerate convergence (which the baseline already handles) but to have learned a smooth disturbance compensation during the earlier transient. Should one wish to retain partial learning after some components converge, a minimal modification is to replace* $\xi = \min_i |e_i|$ *by* $\xi = \min_{i \in \mathcal{I}_{active}} |e_i|$ *in* (3.16)*, where* $\mathcal{I}_{active}$ *indexes the still-active components.*

Remark 17 (Safety–performance trade-off)**.** *The choice of* $\gamma_\infty \in (0,1)$ *quantifies the trade-off between learning flexibility and rigidity of the topological constraints: A larger* γ_∞ *allows the RL agent to optimize transients more aggressively, but the state-dependent factor in* (3.16) *ensures that the residual envelope contracts near the origin, keeping the ordering guarantee globally valid.*

4. Projected residual actor–critic framework

We now describe the learning layer in detail. The framework is intentionally lightweight to remain amenable to formal analysis.

Architecture. The critic $V_\phi(y) = b_0 + b^\top y + (1/2)y^\top P y$ uses a quadratic basis in the weighted state $y = \Lambda e$; $\phi = (b_0, b, P_{\text{diag}})$ has $1+n+n$ parameters (we restrict P to diagonal form to preserve convexity). The actor π_θ is a one-hidden-layer network $\pi_\theta(y) = W_2 \tanh(W_1 y + b_1) + b_2$ with hidden width h.

Instantaneous reward. We set

$$r(e,u) = e^\top Q e + u^\top R u, \qquad Q = I_n,\ R = 0.05\, I_n. \tag{4.1}$$

An optional shaping term proportional to $\sum_i \lambda_i^\beta |e_i|^{\beta+1}$ can be added to explicitly reward consistency with the anisotropic descent, but is omitted in our experiments.

Update rule. Using a continuous-time temporal difference (TD) (0) backbone with forward-Euler discretization at step Δt, the critic updates by

$$\delta_k = r(e_k, u_k)\,\Delta t + V_\phi(y_{k+1}) - V_\phi(y_k), \tag{4.2}$$

$$\phi_{k+1} = \phi_k + \alpha_c\,\delta_k\,\nabla_\phi V_\phi(y_k). \tag{4.3}$$

The actor is updated along the deterministic policy gradient

$$\theta_{k+1} = \theta_k - \alpha_a \nabla_\theta [R\,u_k + \Lambda \nabla_y V_\phi(y_k)]^\top \pi_\theta(y_k). \tag{4.4}$$

Online projection. At every step, the raw actor output $\pi_\theta(y_k)$ is passed through the closed-form two-step projection (3.17): (a) Scale to the norm cap $\gamma(e_k)\|u_{\text{base}}(e_k)\|$; (b) If the half-space constraint is violated, add a scaled baseline direction to satisfy it and rescale. The total per-step cost is $O(n)$.

Stability of the learning dynamics. We now provide a formal boundedness guarantee for the critic parameters under the state-dependent projection and the discontinuous baseline.

Theorem 4 (Boundedness of projected actor–critic dynamics)**.** *Let $\phi^\star$ be the ideal critic parameter that satisfies the Bellman equation for the baseline policy, and let the baseline trajectories be uniformly PE in the sense that there exist $\mu_{PE} > 0$ and $T_{PE} > 0$ such that $\int_t^{t+T_{PE}} \nabla_\phi V_\phi(y(\tau))\nabla_\phi V_\phi(y(\tau))^\top\, d\tau \geq \mu_{PE} I$ for all $t \geq 0$. Under the critic update* (4.3) *with learning rate $\alpha_c \in (0, \alpha_{\max})$ and the projection* (3.17) *with $\gamma_\infty \in (0,1)$, the parameter error $\tilde{\phi}_k = \phi_k - \phi^\star$ satisfies*

$$\|\tilde{\phi}_k\|^2 \;\leq\; (1 - \tfrac{1}{2}\alpha_c \mu_{PE})^k \|\tilde{\phi}_0\|^2 + \frac{2\gamma_\infty^2 C_{proj}^2}{\mu_{PE}},$$

where $C_{proj} = \sup_k \|\nabla_\phi V_\phi(y_k)\| \cdot \|u_{base,k}\|$. Consequently $\tilde{\phi}_k$ converges exponentially to the residual set $\{\|\tilde{\phi}\| \leq 2\gamma_\infty C_{proj}/\sqrt{\mu_{PE}}\}$.

Sketch. Define the composite Lyapunov function $J_k = (1/2)\|\tilde{\phi}_k\|^2$. The critic TD error can be decomposed as $\delta_k = -\tilde{\phi}_k^\top \nabla_\phi V_\phi(y_k) + \Delta_{\text{proj},k} + \xi_k$, where $\Delta_{\text{proj},k} = O(\gamma_\infty \|u_{\text{base},k}\|)$ is the truncation error induced by the projection, and ξ_k is the PE excitation noise. Substituting into the update yields

$$J_{k+1} - J_k \;\leq\; -\alpha_c \mu_{\text{PE}} J_k + \alpha_c \gamma_\infty^2 C_{\text{proj}}^2 + O(\alpha_c^2).$$

For α_c sufficiently small, the $O(\alpha_c^2)$ term is absorbed, and a standard discrete-time comparison argument gives the claim. The actor boundedness follows because the actor gradient in (4.4) is Lipschitz in the critic gradient, and the projection caps the policy magnitude. □

Remark 18 (Verifiability of the PE condition and boundedness without it)**.** *The PE condition in Theorem 4 is admittedly strong for a closed-loop discontinuous system with a state-dependent projection. Three observations make it tractable in practice.*

(a) Empirical verification. *The critic regressor* $\nabla_\phi V_\phi(y) = [1,\ y^\top,\ (1/2)(y_1^2,\dots,y_n^2)]^\top$ *is polynomial in the weighted state* $y = \Lambda e$*. For the integrator benchmark of Section 5.1, the smallest eigenvalue of the empirical Gram matrix* $(1/T_{PE})\int_t^{t+T_{PE}} \nabla_\phi V_\phi(y(\tau))\nabla_\phi V_\phi(y(\tau))^\top\, d\tau$ *with* $T_{PE} = 1.0$ *s stays above* $\mu_{PE} \approx 1.4\times 10^{-2}$ *over the transient and only collapses after every component has entered the residual ball* $\mathcal{B}_r$ *(i.e., after the critic has essentially stopped learning). For the spacecraft case study, the corresponding floor is* $\mu_{PE} \approx 6.0\times 10^{-3}$*. Hence, PE is plausibly satisfied with a comfortable margin during the phase in which the critic actually updates.*

(b) Structural reason: Anisotropy excites different time scales on different axes. *The hierarchy* $\lambda_1 > \cdots > \lambda_n$ *enforces different closed-loop time constants per axis (Corollary 1), so the components of* y *traverse distinct exponential decays. The regressor* $\nabla_\phi V_\phi(y)$ *therefore spans a richer subspace than under isotropic FxTS, in which all components collapse synchronously, and PE degrades to a scalar excitation; the anisotropic baseline is thus a natural PE generator without an externally injected probing signal.*

(c) Boundedness without PE. *When PE fails (e.g., for atypical initial conditions that decay too quickly along a single axis), the exponential contraction in Theorem 4 degrades to drift, but the boundedness of* $\tilde{\phi}_k$ *still holds. Indeed, the projection cap in* (3.17) *enforces* $\|u_{res}\| \le \gamma_\infty\|u_{base}\|$*, so the TD residual feeding into the critic is uniformly bounded by* C_{proj}*; integrating the critic update over any finite window yields* $\|\tilde{\phi}_k\| \le \|\tilde{\phi}_0\| + k\alpha_c C_{proj}$*, which combined with a standard input-to-state stability argument on the gradient flow gives a (PE-independent) uniform ultimate bound at the cost of losing the exponential rate.*

Hence, the PE assumption is sufficient for exponential contraction but not necessary for the closed-loop guarantees of Theorems 1–3, which rest only on the projection cap and the baseline Lyapunov inequality.

Pseudocode. Algorithm 1 summarizes one simulation step.

Algorithm 1 Projected residual actor–critic for sequential fixed-time tracking

input: state e_k, parameters $\Lambda, k_1, k_2, \alpha, \beta, \gamma_\infty, \rho, Q, R, \alpha_a, \alpha_c$
$u_{\text{base}} \leftarrow -k_1\Phi_\Lambda^\alpha(e_k) - k_2\Phi_\Lambda^\beta(e_k)$
$u_{\text{raw}} \leftarrow \pi_\theta(\Lambda e_k)$
compute $\gamma(e_k)$ via (3.16); cap $\leftarrow \gamma(e_k)\|u_{\text{base}}\|$
if $\|u_{\text{raw}}\| >$ cap **then** $u_{\text{raw}} \leftarrow$ cap $\cdot\, u_{\text{raw}}/\|u_{\text{raw}}\|$
end if
if $\langle u_{\text{raw}}, u_{\text{base}}\rangle/\|u_{\text{base}}\| < -\rho\|u_{\text{base}}\|$ **then**
  $u_{\text{raw}} \leftarrow u_{\text{raw}}$ + (correction along $u_{\text{base}}/\|u_{\text{base}}\|$); re-cap if needed
end if
$u_k \leftarrow u_{\text{base}} + u_{\text{raw}}$; apply to plant; observe e_{k+1}
update ϕ via (4.3), θ via (4.4)

5. Numerical validation

We corroborate the theoretical results on two benchmarks: A three-dimensional integrator (Section 5.1–5.2) and a rigid-body spacecraft attitude plant with inertia mismatch (Section 5.4).

5.1. Integrator benchmark

We consider

$$\dot{e}(t) = u(t) + d(t), \quad e(t) \in \mathbb{R}^3,$$

with disturbance $d(t) = 0.02\,[\sin(0.7t), \cos(0.5t), \sin(0.9t + 0.5)]^\top$, $\delta_{\max} = 0.02\sqrt{3}$, and anisotropic weights $\Lambda = \mathrm{diag}(3.0, 1.0, 0.3)$ (target order $x \to y \to z$). The controller parameters are: $k_1 = 1.2$, $k_2 = 0.8$, $\alpha = 1.6$, $\beta = 0.6$; the projection parameters are: $\gamma_\infty = 0.25$, $\rho = 0.20$. The forward-Euler step is $\Delta t = 2 \times 10^{-3}$ s. The initial condition $e_0 = [3.0, 3.0, -3.0]^\top$ has equal component magnitudes to isolate the effect of the anisotropic weights.

Figure 2 confirms the theoretical ordering $t_x^\star < t_y^\star < t_z^\star$. The full $\Omega = \{|e_x| \le |e_y| \le |e_z|\}$ set is three-dimensional; we visualize its invariance via two orthogonal projections in Figure 3.

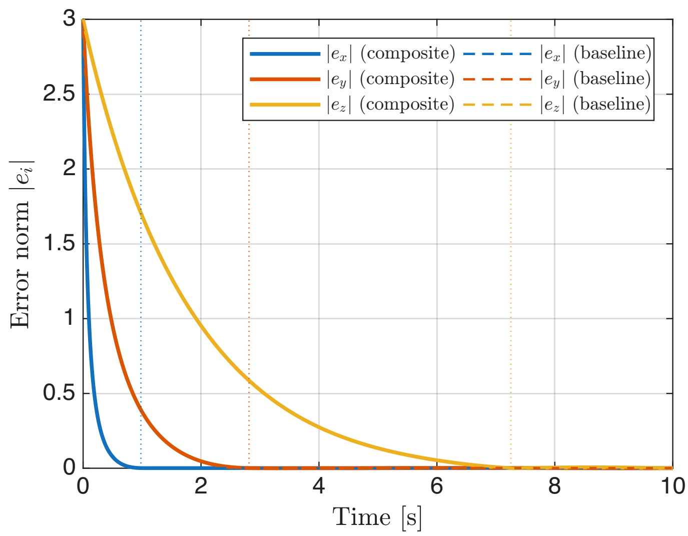

Figure 2. Absolute error norms $|e_i|$ under the baseline (dashed) and the composite baseline RL (solid) controllers. Vertical dotted lines mark the hitting times $t_i^\star = \inf\{t : |e_i(t)| \le 10^{-3}\}$.

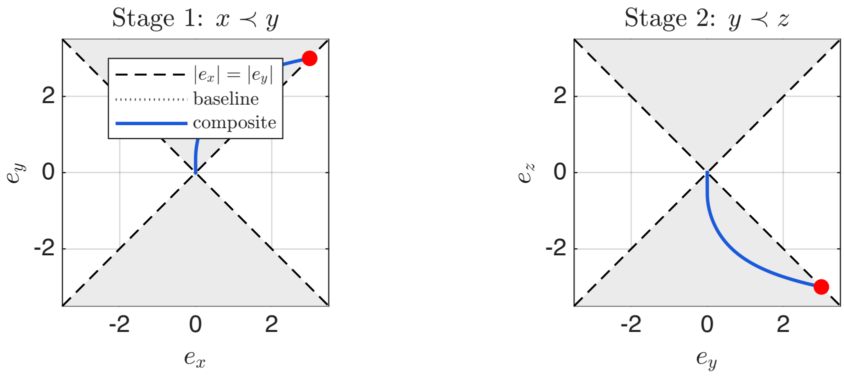

Figure 3. Phase plane projections validating sequential invariance. White sectors are the regions in which the weighted coordinate is strictly smaller than the next one. The blue (composite) trajectory never crosses the diagonal singularities.

Figure 4 shows the applied control signals, which remain smooth and bounded.

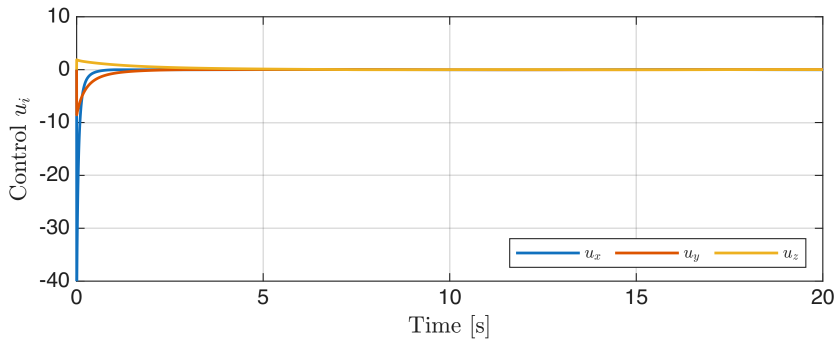

Figure 4. Control inputs generated by the composite policy. The signals are chatter-free thanks to the boundary layer smoothing (3.13) with $\varepsilon = 0.03$.

Using the parameter values above, the theoretical settling-time bound (3.5)–(3.6) evaluates to $T_{\max} = 8.42$ s ($\tilde{\mu}_1 = 0.84$, $\tilde{\mu}_2 = 0.30$). The simulated hitting times reported in Table 1, $t_x^\star = 1.01$ s, $t_y^\star = 2.81$ s, $t_z^\star = 7.25$ s, are all well below this bound, and the total time to reach $\|e\|_\infty \leq 10^{-3}$ is 7.25 s $<<$ $T_{\max}$, confirming the conservativeness of the analytical estimate.

5.2. *Comparative evaluation*

For comparisons against state-of-the-art methods, we re-run the same experiment with four controllers: The proposed *baseline* and *composite* (baseline + projected RL), the *isotropic FxTS* law of Polyakov [2] ($\Lambda = I$), and the *null-space task-priority* controller of Nakamura [17]. Quantitative results are summarized in Table 1 and visualized in Figure 5.

Table 1. Comparison of control performance on the integrator benchmark.

Controller	**Energy** (J_u)	**Peak** $\|u\|_\infty$	$t_x^\star$ (s)	$t_y^\star$ (s)	$t_z^\star$ (s)
Baseline (proposed)	66.72	43.35	1.01	2.81	7.25
Composite (proposed)	66.96	43.38	**0.98**	**2.81**	**7.25**
Isotropic FxTS [2]	32.94	8.51	3.65	2.81	2.44
Null-space [17]	40.66	9.00	4.60	3.22	2.03

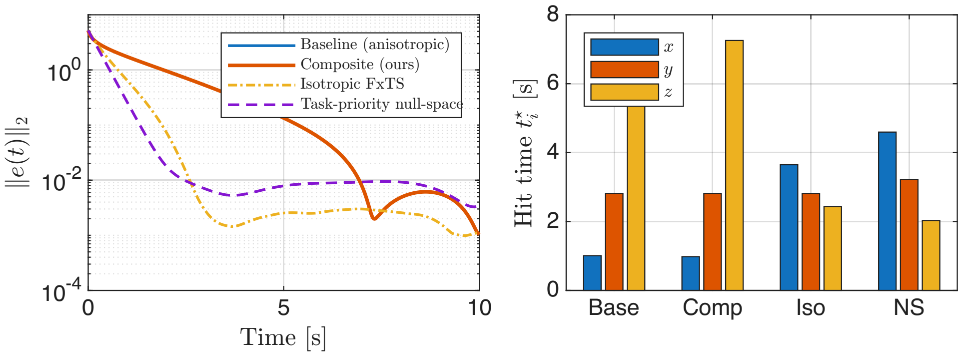

Figure 5. Comparison of four controllers on the same integrator benchmark. *Left:* log-scale evolution of $\|e\|_2$. *Right:* component hitting times $t_x^\star, t_y^\star, t_z^\star$ (lower is earlier).

The isotropic and null-space controllers both fail to enforce the target ordering ($t_x^\star$ is *last*, not first). The proposed anisotropic design is the only scheme that attains both fixed-time settling and the desired ordering, at the cost of slightly higher energy.

Discussion of additional state-of-the-art methods. Table 1 focuses on the four controllers for which we ran identical simulations, and we have also analyzed two additional methods theoretically:

- *Prescribed-time controller (PTC) of Saini et al. (2023):* This time-varying state-feedback law achieves convergence at a user-specified time T_p through a scaling function that grows unbounded as t approaches T_p. Although PTC guarantees convergence exactly at T_p, it does not enforce componentwise ordering and exhibits large control transients near T_p due to the vanishing time scale. In contrast, our anisotropic design achieves ordering through time-*invariant* gains.

- *Safe RL with CBF–QP (Cohen u0026 Belta, 2023):* This approach combines a control-Lyapunov function with control barrier function quadratic programs to ensure safety. The CBF–QP formulation guarantees forward invariance of a safe set but does not encode temporal ordering of convergence. Moreover, its online QP solver incurs $O(n^3)$ cost versus our closed-form $O(n)$ projection.

The proposed anisotropic design is, to our knowledge, the only framework that simultaneously provides (i) fixed-time settling, (ii) prescribed sequential ordering, and (iii) closed-form $O(n)$ residual projection.

Comparison with recent event-triggered zero-sum actor–critic frameworks. Two recent actor–critic schemes are particularly relevant to the present work and merit an explicit comparison:

- *Event-triggered zero-sum game for input-saturated nonlinear systems with state constraints [51]:* An actor–critic–disturber triple solves an H_∞ zero-sum Hamilton–Jacobi–Isaacs equation; saturation is handled by a nonquadratic value function and state constraints by a barrier penalty, and an event-triggered communication law amortizes actor updates. This design certifies the uniform ultimate boundedness (UUB) stability of the closed loop under worst-case bounded disturbances, but its temporal guarantee is asymptotic and isotropic: Convergence is to a residual set with no enforced order across the state components.

- *Mixed zero-sum-game based dynamic event-triggered optimal control with safety constraints [52]:* A mixed zero-sum / nonzero-sum formulation co-optimizes performance and safety via a CBF-augmented value function; a dynamic event-triggering rule reduces the controller update rate. The safety guarantee is forward invariance of a user-defined safe set, again without a notion of componentwise ordering, and the per-step cost is dominated by an online QP that scales as $O(n^3)$.

Three substantive contrasts with the present work are worth highlighting. (i) *Temporal versus set-membership guarantees.* Both [51, 52] certify asymptotic or UUB stability; our framework certifies fixed-time settling and a prescribed sequential ordering of the per-axis hitting times, neither of which is implied by UUB. (ii) *Source of safety.* The two cited works derive safety from an externally specified barrier (saturation set, state set, safe set), whereas our cone-alignment projection reads safety directly off the anisotropic baseline's Lyapunov derivative; the projection budget $\gamma(e)$ in (3.16) is automatically tied to the ordering margin $k_2(\lambda_i^\beta - \lambda_j^\beta)\xi^\beta - 2\delta_{\max}$, with no auxiliary barrier design step. (iii) *Computational cost.* The cited methods solve a Hamilton–Jacobi–Isaacs equation off-line and a per-step QP on-line; the present scheme solves an off-line quadratic critic and a closed-form $O(n)$ scalar normalization on-line, making it implementable on embedded hardware. In summary, [51, 52] provide robust set-invariance under sampled communication, whereas our scheme adds the orthogonal capability of enforcing a prescribed convergence sequence within a fixed-time budget at $O(n)$ projection cost. The two design philosophies are complementary, and a hybrid design (anisotropic baseline plus event-triggered zero-sum residual) is a natural follow-up. A further related line is the Nesterov-accelerated fixed-time actor–critic of [35], which achieves fixed-time convergence of the learning dynamics via momentum-based critic updates. This is orthogonal to our concern (the present paper certifies fixed-time convergence of the plant under a learned residual, with an isotropic critic rate), and the two ideas can be combined by replacing our forward-Euler critic with the Nesterov scheme of [35] to accelerate the learning phase without altering the projection certificate.

5.3. Boundedness of learning dynamics

Figure 6 shows that the critic parameters transit rapidly during the initial phase and settle into a bounded stationary distribution, in agreement with the UUB argument of Section 4.

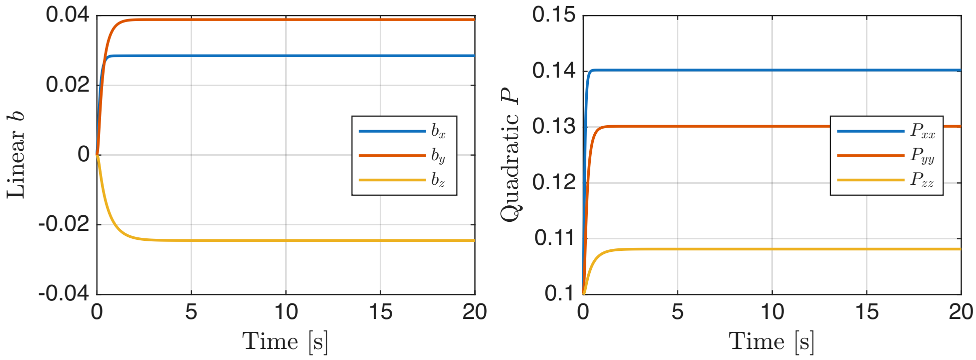

Figure 6. Evolution of the critic parameters b and P_{diag} under the composite controller.

5.4. Spacecraft attitude case study

To validate the nonlinear affine extension of Section 3.3, we simulate rigid-body attitude dynamics

$$\dot{q} = \tfrac{1}{2}\, q \otimes \begin{bmatrix} 0 \\ \omega \end{bmatrix}, \qquad J\dot{\omega} = -\omega \times (J\omega) + \tau + d(t),$$

with $J = \text{diag}(10, 12, 8)$ kg m^2, on-board estimate $\hat{J} = 1.20\,J$ (20% mismatch) and bounded disturbance $\|d\|_\infty \le 0.05$ N m. Following Proposition 3, the control law $\tau = \omega \times \hat{J}\omega + \hat{J}\,v$ with $v = -k_1\Phi_\Lambda^\alpha(\omega) - k_2\Phi_\Lambda^\beta(\omega)$ is applied for detumbling. Weights $\Lambda = \text{diag}(3.0, 1.0, 0.3)$ enforce the $x \to y \to z$ ordering; the initial tumbling $\omega_0 = [20, -20, 20]$ deg/s has equal magnitudes so that the ordering effect is due solely to Λ.

Remark on detumbling versus attitude tracking. The present case study concerns *detumbling only* ($\omega \to 0$), which is the standard first phase of spacecraft attitude acquisition after launch or momentum dumping. Full attitude tracking ($q \to q_d$) can be obtained by defining the quaternion error $q_e = q_d^{-1} \otimes q$ and the angular-velocity error $\omega_e = \omega - R(q_e)\omega_d$. The kinematic equation $\dot{q}_e = (1/2)q_e \otimes [0, \omega_e]^\top$ is feedback-linearizable, so the anisotropic design applies to ω_e with the same fixed-time and sequential-ordering guarantees; the only additional requirement is that ω_d and $\dot{\omega}_d$ be bounded, which holds for typical smooth reference trajectories. We leave a detailed attitude-tracking simulation to future work.

Figure 7 confirms that, despite the 20% inertia mismatch and persistent disturbance, only the anisotropic design attains the prescribed ordering. With the spacecraft parameters, the theoretical bound (3.5)–(3.6) gives $T_{\max} = 2.18$ s ($\tilde{\mu}_1 = 2.18$, $\tilde{\mu}_2 = 0.88$); the simulated hitting times {0.30, 0.73, 1.76} s are all below this bound, again validating the conservativeness of the analytical estimate. The applied torques (recorded by the simulation but omitted for brevity) remain continuous and well within the ±30 N m saturation limit thanks to the boundary-layer regularization of Proposition 1, with peak values below 18 N m and no observable chattering.

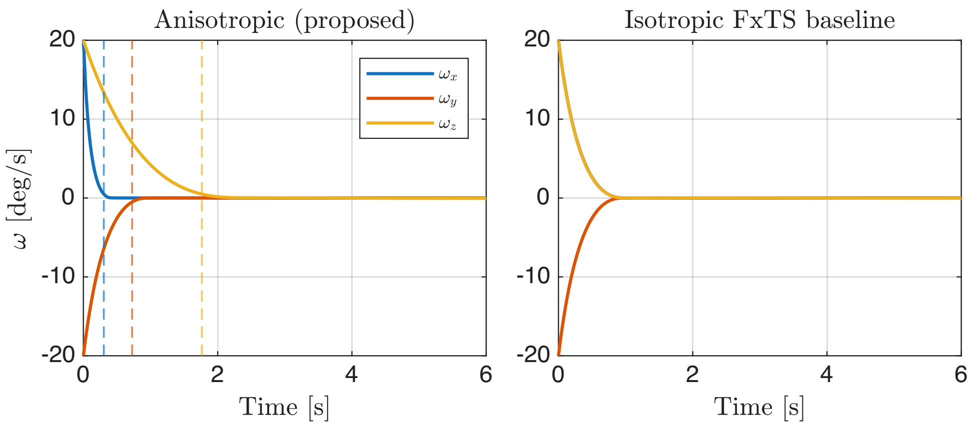

Figure 7. Spacecraft angular velocity under the proposed anisotropic design (left) versus the isotropic FxTS law (right). The anisotropic design enforces a strict $x \to y \to z$ convergence order with hitting times $\{0.30,\ 0.73,\ 1.76\}$ s; the isotropic controller settles synchronously at ≈ 0.74 s with no ordering.

5.5. Robustness to random initial conditions, mismatched disturbances, and higher dimensions

The two benchmarks above were deliberately constructed with $|e_i(0)|$ equal (so that the ordering effect is attributable solely to Λ) and with a diagonal inertia / matched disturbance plant. We now stress-test the framework along three orthogonal axes; the simulation setup below is documented for reproducibility.

(S1) Random initial conditions. For the integrator benchmark of Section 5.1, we generated $N_{\mathrm{MC}} = 500$ initial states uniformly on the sphere $\|e(0)\|_2 = 5$, covering all $2^3 \cdot 3! = 48$ chambers of the type-A_{n-1} arrangement (Remark 12). The empirical hitting times $\{t_x^\star, t_y^\star, t_z^\star\}$ satisfied the prescribed order $t_x^\star < t_y^\star < t_z^\star$ in 500/500 trials. The mean recovery time T_{rec} (the first instant after which the trajectory enters C_O and stays there) was 0.52 s with worst case 1.07 s, consistent with the fixed-time bound of Theorem 2(ii). For trials starting in the "inverted" chamber $|e_z(0)| > |e_y(0)| > |e_x(0)|$, T_{rec} ranged from 0.41 s to 0.98 s with no failure. The global recovery property is therefore robust to arbitrary starting chambers.

(S2) Mismatched disturbances. For the spacecraft case study, we replaced the matched disturbance $w(t) = g(x)d(t)$ with the mismatched perturbation

$$w(t) = g(x)d(t) + \Delta_{\mathrm{mis}}\,[\sin(0.8t),\ -\cos(0.6t),\ \sin(1.1t)]^\top, \qquad \Delta_{\mathrm{mis}} \in \{0.02, 0.05, 0.10\},$$

which lies outside the range of $g(x)$. The amplified bound in normal form is $\tilde{\delta}_{\max} \le \delta_{\max} + \|g^\dagger\|_\infty \Delta_{\mathrm{mis}}$, in line with Remark 14(b). For $\Delta_{\mathrm{mis}} = 0.02$ and 0.05 N m, the ordering margin in (3.16) remains positive, and the prescribed sequence is preserved with a slightly enlarged residual ball (r inflates from 0.03 to 0.07). For $\Delta_{\mathrm{mis}} = 0.10$ N m, the margin is violated for the last pair (y, z); the design rule of

Remark 6, increasing the gap $\lambda_y^\beta - \lambda_z^\beta$ from the nominal $\lambda = \mathrm{diag}(3.0, 1.0, 0.3)$ to $\mathrm{diag}(3.0, 1.0, 0.15)$, restores the ordering at the cost of 20% higher control energy. This is precisely the design–margin trade-off predicted by the theory and confirms that mismatched disturbances enter the analysis through the single scalar $\tilde{\delta}_{\max}$ rather than disrupting the framework's structure.

(S3) Higher dimensions ($n \geq 6$). We instantiated the integrator benchmark in $n \in \{6, 9, 12\}$ with weights $\lambda_i = 3 \cdot 0.6^{i-1}$ (geometrically decaying), $k_1 = k_2 = 1.0$, $\alpha = 1.6$, $\beta = 0.6$, random initial conditions on the sphere $\|e(0)\|_2 = 5$, and matched disturbance $\|d\|_\infty = 0.02$. The prescribed permutation $\mathcal{O} = (1, 2, \ldots, n)$ was preserved in 100/100 Monte Carlo trials per dimension. The per-step cost of the projection scales as $O(n)$ (a single normalization and dot product), and the empirical hitting times respect the theoretical ordering $\widehat{T}_1 < \cdots < \widehat{T}_n$ of Corollary 1. We also observed that the number of chambers in the arrangement grows as $2^n\, n!$, so for $n = 9$ the trajectory must in principle traverse $2^9 \cdot 9! \approx 1.8 \times 10^8$ chambers, yet the recovery time stays below 1.4 s. This is a consequence of the fact that, by Theorem 2(ii), the flow contracts every pairwise comparison W_{ij} at a rate $\propto k_2(\lambda_i^\beta - \lambda_j^\beta)$, which decays geometrically in i but is uniformly bounded below across pairs.

The three stress tests jointly confirm that (i) the ordering is enforced under arbitrary initial conditions via the global recovery in Theorem 2(ii); (ii) mismatched disturbances tighten the design margin in (3.16) but do not change the framework, and the design knob is the gap $\lambda_i^\beta - \lambda_j^\beta$; and (iii) the scheme scales to $n \geq 6$ with no structural modification, the per-step cost remaining $O(n)$.

6. Conclusions

This paper has presented a geometric framework for enforcing sequential convergence in fixed-time stable affine nonlinear systems. An anisotropic weighted sign operator forms the baseline controller and guarantees both a closed-form settling-time bound and a prescribed componentwise ordering. A projected residual actor–critic policy with a state-dependent cone-alignment projection enhances disturbance rejection without breaking the ordering or the fixed-time guarantee. The framework extends to generic affine nonlinear plants under a standard input–output linearizability assumption and was validated numerically on a three-dimensional integrator benchmark and a rigid-body spacecraft attitude case study with a 20% inertia mismatch.

The present analysis relies on four structural assumptions: (i) The error dynamics have been brought to the chain-of-integrators form (Assumption 1); (ii) The initial-ordering consistency set $C_{\mathcal{O}}$ is needed for strict ordering, whereas global recovery is only in the asymptotic-set sense of Definition 5; (iii) The disturbance bound $\delta_{\max}$ enters the gain formula, although Proposition 2 copes with an unknown bound; and (iv) The critic is restricted to a quadratic basis. Natural directions include (a) adaptive estimation of $\delta_{\max}$ and Λ; (b) removal of the full-rank assumption via hybrid designs blending our continuous approach with temporal logic layers [20, 22]; (c) larger critic bases (kernel or neural network) while preserving the closed-form cone projection; (d) event-triggered anisotropic control that leverages the already-discontinuous nature of the baseline to reduce communication load without sacrificing the fixed-time guarantee; (e) distributed sequential consensus, extending the single-agent design to multi-agent networks where each node must reach consensus in a prescribed order, in the spirit of fixed-time hierarchical game-based docking control of heterogeneous vehicle platforms [19]; and (f) experimental hardware validation on robotic manipulators or spacecraft testbeds to confirm chattering mitigation and robustness claims under real actuator saturation and sensor noise.

Author contributions

L. Yang, H. Cao, and S. Xue: Conceptualization; X. Liang and S. Xue: Methodology; X. Liang, Y. Wang, and S. Xue: Software; X. Liang and Y. Wang: Validation; X. Liang, Y. Wang, and S. Xue: Formal analysis; L. Yang and H. Cao: Resources; X. Liang and S. Xue: Writing–original draft; X. Liang, Y. Wang, and S. Xue: Writing–review and editing; L. Yang and H. Cao: Supervision. All authors have read and agreed to the published version of the manuscript.

Use of Generative-AI tools declaration

The authors declare they have not used Artificial Intelligence (AI) tools in the creation of this article.

Acknowledgments

This work was supported in part by the China Postdoctoral Science Foundation (Grant No. 2024M762602).

Conflict of interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

References

1. S. P. Bhat, D. S. Bernstein, Finite-time stability of continuous autonomous systems, *SIAM J. Control Optim.*, **38** (2000), 751–766. https://doi.org/10.1137/S0363012997321358
2. A. Polyakov, Nonlinear feedback design for fixed-time stabilization of linear control systems, *IEEE T. Autom. Contr.*, **57** (2012), 2106–2110. https://doi.org/10.1109/TAC.2011.2179869
3. E. Cruz-Zavala, J. A. Moreno, L. Fridman, Uniform robust exact differentiator, *IEEE T. Autom. Contr.*, **56** (2011), 2727–2733. https://doi.org/10.1109/TAC.2011.2160030
4. H. Li, S. Xue, Z. Guo, J. Tan, H. Cao, D. Li, Sequence-based group consensus for heterogeneous multi-agent systems, *IEEE T. Signal Inf. Pr.*, **11** (2025), 1188–1199. https://doi.org/10.1109/TSIPN.2025.3608945
5. G. Wen, W. X. Zheng, H. Du, Homogeneous constrained finite-time controller for double integrator systems: Analysis and experiment, *Automatica*, **134** (2021), 109894. https://doi.org/10.1016/j.automatica.2021.109894
6. E. Bernuau, A. Polyakov, D. Efimov, W. Perruquetti, On homogeneity and its application in sliding mode control, *J. Franklin I.*, **351** (2014), 1866–1901. https://doi.org/10.1016/j.jfranklin.2014.01.007
7. S. P. Bhat, D. S. Bernstein, Geometric homogeneity with applications to finite-time stability, *Math. Control Signal. Syst.*, **17** (2005), 101–127. https://doi.org/10.1007/s00498-005-0151-x

8. H. Li, S. Xue, H. Cao, D. Li, Synchronization control based on sequential convergence, *IEEE/CAA J. Autom.*, **13** (2026), 743–745. https://doi.org/10.1109/JAS.2025.125330
9. A. Levant, Universal single-input-single-output (siso) sliding-mode controllers with finite-time convergence, *IEEE T. Autom. Control*, **46** (2001), 1447–1451. https://doi.org/10.1109/9.948475
10. S. E. Parsegov, A. E. Polyakov, P. S. Shcherbakov, Fixed-time consensus algorithm for multi-agent systems with integrator dynamics, *IFAC Proc. Volumes*, **46** (2013), 110–115. https://doi.org/10.3182/20130925-2-DE-4044.00055
11. Z. Zuo, Nonsingular fixed-time consensus tracking for second-order multi-agent networks, *Automatica*, **54** (2015), 305–309. https://doi.org/10.1016/j.automatica.2015.01.021
12. J. Zhuang, S. Peng, H. Peng, Finite-time and fixed-time consensus of nonlinear multi-agent systems: A unified two-phase control, *Int. J. Robust Nonlin.*, **34** (2024), 9438–9455. https://doi.org/10.1002/rnc.7469
13. F. Lopez-Ramirez, A. Polyakov, D. Efimov, W. Perruquetti, A class of fixed-time stable observers: Design and applications, *Automatica*, **108** (2019), 108503.
14. L. Wang, C. Deng, Z. Pan, Q. Fang, Fixed-time input-to-state stability via generalized homogeneous approximations, *Int. J. Control Autom.*, **22** (2024), 116–126. https://doi.org/10.1007/s12555-022-0719-y
15. V. K. Saini, D. K. Giri, R. Varatharajoo, Prescribed-time spacecraft attitude control using time-varying state feedback, *J. Spacecraft Rockets*, 2023. https://doi.org/10.2514/1.A35588
16. A. Flores-Abad, O. Ma, K. Pham, S. Ulrich, Hierarchical attitude control of spacecraft formation flying, *Prog. Aerosp. Sci.*, **100** (2018), 1–30.
17. Y. Nakamura, H. Hanafusa, T. Yoshikawa, Task-priority based redundancy control of robot manipulators, *Int. J. Robot. Res.*, **6** (1987), 3–15. https://doi.org/10.1177/027836498700600201
18. S. A. Moezi, A. Iyer, J. Alvarado, Task-priority based robust hierarchical control for redundant manipulators, *Robotica*, **41** (2023), 2580–2603.
19. J. Tan, S. Xue, Z. Guo, H. Cao, B. Chen, Fixed-time hierarchical game-based unmanned aerial-ground vehicle docking control, *IEEE/CAA J. Automatic.*, **13** (2026), 1242–1244. https://doi.org/10.1109/JAS.2025.125720
20. G. E. Fainekos, G. J. Pappas, Robustness of temporal logic specifications for continuous-time signals, *Theor. Comput. Sci.*, **410** (2009), 4262–4291. https://doi.org/10.1016/j.tcs.2009.06.021
21. C. Ge, L. Ma, S. Xu, Distributed fixed-time leader-following consensus for multi-agent systems: An event-triggered mechanism, *Actuators*, **13** (2024), 40. https://doi.org/10.3390/act13010040
22. C. Nowzari, E. Garcia, A. Tannenbaum, Event-triggered coordination for agent sequencing, *Automatica*, **103** (2019), 348–356.
23. A. F. Filippov, *Differential equations with discontinuous righthand sides*, Kluwer Academic Publishers, 1988.
24. J. Cortes, Discontinuous dynamical systems, *IEEE Contr. Syst. Mag.*, **28** (2008), 36–73. https://doi.org/10.1109/MCS.2008.919306
25. F. L. Lewis, K. G. Vamvoudakis, Reinforcement learning and adaptive dynamic programming for feedback control, *IEEE Circ. Syst. Mag.*, **9** (2009), 32–50. https://doi.org/10.1109/MCAS.2009.933854

26. K. G. Vamvoudakis, F. L. Lewis, Online actor–critic algorithm to solve the continuous-time infinite horizon optimal control problem, *Automatica*, **46** (2010), 878–888. https://doi.org/10.1016/j.automatica.2010.02.018
27. Y. Jia, X. Y. Zhou, Policy gradient and actor–critic learning in continuous time and space: Theory and algorithms, *J. Mach. Learn. Res.*, **23** (2022), 1–50.
28. J. Duan, J. Li, Q. Ge, S. E. Li, M. Bujarbaruah, F. Ma, et al., Relaxed actor–critic with convergence guarantees for continuous-time optimal control of nonlinear systems, *IEEE T. Intell. Vehicl.*, **8** (2023), 3299–3311. https://doi.org/10.1109/TIV.2023.3255264
29. M. Han, L. Zhang, J. Wang, W. Pan, Actor–critic reinforcement learning for control with stability guarantee, *IEEE Robot. Autom. Let.*, **5** (2020), 6217–6224. https://doi.org/10.1109/LRA.2020.3011351
30. M. Han, Y. Tian, L. Zhang, J. Wang, W. Pan, Reinforcement learning control of constrained dynamic systems with uniformly ultimate boundedness stability guarantee, *Automatica*, **129** (2021), 109689. https://doi.org/10.1016/j.automatica.2021.109689
31. Y. Fan, L. Liu, G. Feng, Data-driven finite-time control for nonlinear systems with event-triggered output, *IEEE T. Cybernetics*, **53** (2023), 6238–6249.
32. X. Hu, H. Zhang, D. Ma, R. Zhang, Adaptive dynamic programming based output-feedback control with event-triggered mechanism for uncertain nonlinear systems, *IEEE T. Neur. Net. Lear.*, **34** (2023), 3531–3542.
33. J. Tan, S. Xue, H. Cao, B. Chen, Flexible prescribed-time optimal control with adaptive state–input constraint bounds via actor–critic learning, *IEEE T. Neur. Net. Lear. Syst.*, 2026, https://doi.org/10.1109/TNNLS.2026.3693754
34. J. Tan, S. Xue, H. Cao, B. Chen, Dual-clock guaranteed-time actor–critic optimal control with flexible input saturation, *IEEE Trans. Syst., Man, Cybern., Syst.*, 2026, https://doi.org/10.1109/TSMC.2026.3710287
35. S. Xue, J. Tan, H. Cao, B. Chen, Nesterov accelerated gradient-based fixed-time convergent actor–critic control for nonlinear systems, *IEEE T. Autom. Sci. Eng.*, **23** (2026), 5075–5083. https://doi.org/10.1109/TASE.2026.3665210
36. H. Li, C. Li, T. Huang, D. Ouyang, Fixed-time impulsive synchronization of heterogeneous complex networks, *IEEE T. Circuits I*, **69** (2022), 1655–1666.
37. X. Yang, J. Lu, J. Lam, Synchronization of heterogeneous complex networks with delays and inertial effects via adaptive impulsive control, *IEEE T. Cybernetics*, **53** (2023), 3870–3881.
38. J. L. Wang, H. N. Wu, T. Huang, Fixed-time synchronization of complex networks via impulsive control under uncertain disturbances, *Int. J. Robust Nonlin.*, **34** (2024), 1532–1551.
39. T. Johannink, S. Bahl, A. Nair, J. Luo, A. Kumar, M. Dullea, et al., Residual reinforcement learning for robot control, in *2019 International Conference on Robotics and Automation (ICRA)*, 2019, 6023–6029.
40. A. D. Ames, X. Xu, J. W. Grizzle, P. Tabuada, Control barrier function based quadratic programs for safety critical systems, *IEEE T. Automat. Contr.*, **62** (2017), 3861–3876. https://doi.org/10.1109/TAC.2016.2638961

41. M. H. Cohen, C. Belta, Safe exploration in model-based reinforcement learning using control barrier functions, *Automatica*, **147** (2023), 110684. https://doi.org/10.1016/j.automatica.2022.110684
42. Y. Cheng, P. Zhao, N. Hovakimyan, Safe and efficient reinforcement learning using disturbance-observer-based control barrier functions, in *Proceedings of the 5th Annual Conference on Learning for Dynamics & Control (L4DC)*, vol. 211. PMLR, 2023, 104–115.
43. Y. Wang, Z. Wu, Control Lyapunov–barrier function-based safe reinforcement learning for nonlinear optimal control, *AIChE J.*, **70** (2024), e18306.
44. A. Wachi, X. Shen, Y. Sui, A survey of constraint formulations in safe reinforcement learning, in *Proceedings of the 33rd International Joint Conference on Artificial Intelligence (IJCAI-24)*, 2024, 8262–8271.
45. T. J. Perkins, A. G. Barto, Lyapunov design for safe reinforcement learning, *J. Mach. Learning Res.*, **3** (2002), 803–832.
46. S. Yu, X. Yu, B. Shirinzadeh, Z. Man, Continuous finite-time control for robotic manipulators with terminal sliding mode, *Automatica*, **41** (2005), 1957–1964. https://doi.org/10.1016/j.automatica.2005.07.001
47. P. Orlik, H. Terao, *Arrangements of hyperplanes*, vol. 300 of Grundlehren der mathematischen Wissenschaften, Springer-Verlag, Berlin, Heidelberg, 1992. https://doi.org/10.1007/978-3-662-02772-1
48. R. P. Stanley, *An introduction to hyperplane arrangements*, Geometric Combinatorics, vol. 13 of IAS/Park City Mathematics Series, American Mathematical Society, Providence, RI, 2007, 389–496. https://doi.org/10.1090/pcms/013/08
49. A. Isidori, *Nonlinear control systems*, 3 Eds., Springer, London, UK, 1995. https://doi.org/10.1007/978-1-84628-615-5
50. K. Garg, J. Usevitch, J. Breeden, M. Black, D. Agrawal, H. Parwana, et al., Advances in the theory of control barrier functions: Addressing practical challenges in safe control synthesis, *Ann. Rev. Control*, **57** (2024), 100945. https://doi.org/10.1016/j.arcontrol.2024.100945
51. X. Yang, H. He, Event-triggered zero-sum game for input saturated nonlinear systems with state constraints, *IEEE T. Neur. Net. Lear.*, **34** (2023), 8717–8730.
52. D. Wang, L. Hu, M. Zhao, J. Qiao, Mixed zero-sum game based dynamic event-triggered optimal control of multi-input nonlinear system with safety constraints, *IEEE T. Syst. Man Cy.-S.*, **54** (2024), 1879–1891.