International Journal of Adaptive Control and Signal Processing

RESEARCH ARTICLE

Robust Adaptive Kolmogorov-Arnold Neural Control

Adilkhan Salkimbayev | Khawaja S. Haider

[1]School of Information Technology and Engineering, Kazakh-British Technical University, Almaty, Kazakhstan |

Correspondence: Adilkhan Salkimbayev (a_salkimbayev@kbtu.kz) | Khawaja S. Haider (s.khawaja@kbtu.kz)

Received: **Revised:** **Accepted:**

Keywords: adaptive control | model predictive control | Kolmogorov-Arnold networks | neuro-symbolic control | real-time control | interpretable neural networks | non-minimum phase systems

ABSTRACT

Real-time control of non-minimum phase systems is limited on embedded platforms by the cost of online optimization and weak guarantees under model mismatch. Adaptive schemes remove the optimization, but their guarantees rest on persistent excitation (PE) — a condition that regulation to a fixed setpoint structurally destroys, restorable only by probing that conflicts with the control objective. We propose a Lipschitz-continuous adaptive architecture using recursive least squares to update Kolmogorov-Arnold networks for interpretable control of control-affine systems with a left-invertible input map, and establish that its guarantee carries no PE requirement: the damping closing the Lyapunov argument comes from a covariance bound the update law enforces at every step, not from excitation of the regressor, so bounded tracking survives loss of excitation. The analysis certifies input-to-state stability to the lumped approximation-and-disturbance error and uniform ultimate boundedness over the entire physically realizable operating envelope. The guarantee is exercised, not merely compatible: on the quadruple-tank benchmark the feature Gram holds numerical rank seven of ten, so excitation is measurably absent. Over a 50-seed Monte-Carlo under adversarial parameter drift, the controller improves on the aggregate tracking of LTV-MPC while matching a radial-basis-function adaptive baseline, at roughly half the per-step cost, on a branch-free evaluation with an a priori worst-case bound, while conceding a 1.6–3.6× actuator-total-variation advantage to the optimizer. Because the distilled basis is frozen offline and only its linear coefficients adapt, the deployed law is always an explicit low-degree polynomial, inspectable from its current coefficients — suiting resource-constrained embedded platforms.

1 | Introduction

Model predictive control (MPC) and robust control strategies have demonstrated exceptional performance in numerical benchmarks, displaying extraordinary resilience and disturbance rejection [1]. These results have extensively validated the theoretical stability guarantees established in foundational works [2, 3]. However, the computational requirements for such optimization-based techniques have impeded their widespread deployment in industrial plants. Despite evolution in quadratic program (QP) solvers [4] and the democratization of convex optimization tools [5, 6], a stalemate has emerged: industrial sectors remain reliant on PID control, largely unable to adopt MPC due to the hardware requirements of online optimization. Thus, a need for a computationally efficient nonlinear controller persists.

Adaptive control is the classical answer to that need, and it exchanges the optimization for a different prerequisite. Guarantees for online function-approximation schemes are conventionally stated under persistent excitation (PE), which is what makes the parameter error converge and the certificate close. Regulation, however, is precisely the operating regime that destroys PE: driving a plant to a fixed setpoint and holding it there collapses the regressor onto a low-dimensional subspace, and the excitation can be restored only by injecting probing signals that work against the control objective. A guarantee contingent on PE is therefore weakest exactly where a process controller spends nearly all of its service life. The central theoretical claim of this work (Corollary 2) is that the proposed architecture needs no such contingency: its closed-loop boundedness certificate contains no PE hypothesis, because the damping that closes the Lyapunov argument is enforced by the update law — through a projected covariance bound — rather than drawn from the trajectory. The benchmark reported here is not a favourable case for this claim but an adverse one: the feature Gram retains numerical rank seven of ten, so excitation is measurably absent and the PE-free argument is load-bearing rather than decorative.

Recent data-driven approaches have sought to replace resource-intensive online optimization with learnable architectures. While methods like SINDYc [7] improve interpretability through sparse identification, the broader adoption of neural control has been stifled by the intractable structure and uninterpretable nature of standard multi-layer perceptrons (MLPs) [8]. By contrast, Kolmogorov-Arnold networks (KANs) provide an alternative: by utilizing the Kolmogorov-Arnold theorem of universal function approximation [9, 10], KANs are capable of providing a tractable symbolic structure, possessing faster scaling laws than MLPs [11]. As a result, KANs enabled research on interpretable neural network-based system identification and control [12, 13].

In previous work [14], the feasibility of deploying static KANs for feedback control was established: the distilled law admits a forensic audit of its own output and is structurally simple enough to inspect term by term. That work also delimits what distillation costs. The symbolic read-out is itself a dominant error source — an off-the-shelf extraction inflates a 4.4% imitation error to 11.0% of the actuator range in the non-minimum-phase regime — and is recovered only by re-estimating the coefficients on the KAN-selected support by convex least squares; on a matched term budget that support is moreover statistically indistinguishable from one chosen by sparse regression. What the symbolic step is credited with there is constant-time execution and auditability, not superior accuracy. The present work inherits that reading rather than a fidelity claim. What it addresses instead is that the earlier process relied on offline training, leaving the controller vulnerable to unmodeled dynamics and parameter drift once deployed.

Online function-approximation adaptive control is a mature paradigm: radial-basis-function (RBF) networks [15] and feedforward neural networks [16, 17] have long been used as in-loop regressors whose weights adapt online, with Lyapunov-based boundedness guarantees rooted in classical robust adaptive control [18, 19, 20, 21]. The distinction between the proposed controller and this paradigm is not adaptation per se but the form of the regressor: where RBF/MLP schemes adapt a black-box basis, the proposed controller adapts the linear weights of an interpretable, symbolically distilled KAN basis using recursive least squares (RLS) [22], retaining the auditability of the static KAN control law while recovering robustness to drift. This positions the contribution between black-box neural-adaptive control and sparse-identification approaches such as SINDYc [7].

In this article, the Lipschitz-continuous adaptive RLS-KAN (LCA-RLS-KAN) is proposed. The contributions of this paper are fourfold: first, RLS is introduced as an update law for the KAN-distilled polynomials that operates online; second, a Lyapunov-based theorem establishing input-to-state stability (ISS) and uniform ultimate boundedness (UUB) of the tracking and weight errors over the entire physically realizable operating envelope, with the (separately, a priori bounded) integral state handled via a leaky integrator confined to an explicit band, is provided, together with a corollary quantifying practical tracking accuracy in the ideal case; third, and centrally, that certificate is shown to be free of any persistent-excitation requirement (Corollary 2) — boundedness is inherited from a covariance bound the update law enforces by projection, so the guarantee holds in the unexcited single-setpoint regime where parameter identification provably fails, and the reported benchmark is verified to be such a regime; fourth, a 50-seed Monte-Carlo study on a high-fidelity quadruple-tank benchmark — conducted against a reachable reference and an LTV-MPC given a horizon that spans the plant's inverse response, so that the baseline is not handicapped — shows that the controller improves on LTV-MPC in aggregate tracking while matching a black-box RBF adaptive baseline, at roughly half the per-step cost of either MPC configuration, on a worst-case-bounded, branch-free evaluation, while conceding a 1.6–3.6× actuator-total-variation advantage to the optimizer. Against the RBF baseline, which shares the identical adaptation harness, control performance is statistically indistinguishable, so the distilled basis is credited with interpretability rather than with tracking or smoothness.

2 | Related Work

Function-approximation adaptive control. The closest methodological relatives adapt an in-loop function approximator online under a Lyapunov boundedness guarantee: RBF networks [15] and feedforward/recurrent neural networks for robotic and nonlinear systems [16, 17]. These supply an uninterpretable regressor whose internal units carry no physical meaning. The proposed method retains the adaptive-update machinery of this lineage but replaces the regressor with an interpretable, symbolically distilled KAN basis; the controlled RBF baseline of Sec. 5.2 is precisely a representative of this class, so that comparison isolates the effect of the basis itself.

Robust and adaptive control foundations. The robustification employed here — exponential forgetting with bounded-gain covariance handling [23], an optional σ-modification (Sec. 4.7), and leakage or anti-windup on the integral state — is classical [18, 19, 20, 21, 24]. Input-to-state stability [25, 26] is the property certified in Sec. 4.6.

Learning-based and approximate MPC. A parallel line replaces online optimization with a learned surrogate of an MPC policy, typically an MLP [8], or identifies sparse governing dynamics for control [7]. Such surrogates inherit the opacity of their underlying networks. Distillation into a symbolic law (Sec. 4.4) instead yields an inspectable controller; the present contribution is to make that distilled law adaptive and to equip it with a stability certificate.

Kolmogorov–Arnold networks and neuro-symbolic control. KANs [11], grounded in the Kolmogorov–Arnold representation theorem [9, 10], have recently been applied to interpretable system identification [12] and constrained control [13]. This work exploits their distillability as the route to a symbolic, online-adaptive control law.

The quadruple-tank benchmark. Johansson's process [27] is a standard multivariable, adjustable-zero, NMP testbed [28], studied under model-based [29] and distributed/decentralized MPC [30]; it is used here as a representative instance of the control-affine class of Sec. 3.1. Table 1 positions these approaches on the axes that matter for embedded deployment.

TABLE 1 | Positioning relative to representative control approaches. ✓: yes; ∼: partial; ×: no.

Approach	Online adapt.	Interp. law	Stability cert.	Determ. per step
LTV / nonlinear MPC	∼	×	✓	×
Approximate MPC (MLP) [8]	×	×	∼	✓
RBF/MLP adaptive [15, 17]	✓	×	✓	✓
SINDYc [7]	×	✓	∼	✓
Static symbolic KAN [14]	×	✓	∼	✓
LCA-RLS-KAN (ours)	✓	✓	✓	✓

TABLE 2 | Notation used in the stability analysis.

Symbol	Meaning
x	plant state (here, tank levels $h_{1...4}$)
$e = x - x_{\mathrm{ref}}$	tracking error
$\hat{\theta},\ \theta^*$	estimated / ideal linear KAN weights
$\tilde{\theta} = \hat{\theta} - \theta^*$	weight estimation error
$\hat{d},\ d^*$	integral state / unknown constant bias
$\tilde{d} = \hat{d} - d^*$	integral-state error (bounded a priori by construction)
$\zeta = [\,e^T\ \tilde{\theta}^T\,]^T$	composite state vector
$\Phi_{\mathrm{KAN}}(x)$	distilled nonlinear regressor vector
$\delta(x)$	lumped bounded disturbance (see text)

3 | Problem Formulation

3.1 | System Class and Notation

The controller is developed for the class of control-affine multivariable systems

$$\dot{x} = f(x) + B\,u, \qquad x \in \mathbb{R}^n,\ \ u \in \mathbb{R}^m, \tag{1}$$

where $f : \Omega \to \mathbb{R}^n$ is continuous on a compact operating set $\Omega \subset \mathbb{R}^n$ and $B \in \mathbb{R}^{n\times m}$ has full column rank (left-invertible), so that the left pseudoinverse $B^+ = (B^T B)^{-1} B^T$ satisfies $B^+ B = I$. The plant may be NMP and may be underactuated ($m < n$); the residual $r_{\mathrm{ua}} := (BB^+ - I)(\cdot)$ arising from underactuation is treated as a bounded disturbance (Lemma 1). The control objective is tracking of a feasible, bounded reference $x_{\mathrm{ref}}(t)$ with bounded $\dot{x}_{\mathrm{ref}}(t)$, with all internal signals remaining bounded under non-vanishing model mismatch. The quadruple-tank process of Sec. 3.2 is a representative instance of (1).

Throughout, $\|\cdot\|$ is the Euclidean norm and $\|\cdot\|_\infty$ the essential-supremum norm; for $M = M^T \succ 0$, $\lambda_{\min}(M), \lambda_{\max}(M)$ denote its extreme eigenvalues and $\|v\|_M^2 := v^T M v$. The error variables used in the analysis are summarized in Table 2. The Lyapunov analysis (Sec. 4.6) is conducted over the composite state $\zeta = [\,e^T\ \tilde{\theta}^T\,]^T$; the integral-state error $\tilde{d}$ is tracked separately and bounded directly by construction (Lemma 2) rather than through ζ.

3.2 | Benchmark Instance: Quadruple-Tank Process

To compare the controllers, Johansson's quadruple-tank process is used — a nonlinear multivariable benchmark consisting of four interconnected liquid tanks and two control inputs, exhibiting strong cross-coupling and NMP behavior [27]. It is an instance of (1) with $n = 4$, $m = 2$, and has since become a standard benchmark for multivariable, decentralized, and distributed control design [29, 30]. The process can be configured to minimum-phase or NMP regimes, enabling a controlled assessment of performance degradation and robustness challenges in adversarial control settings [27, 28].

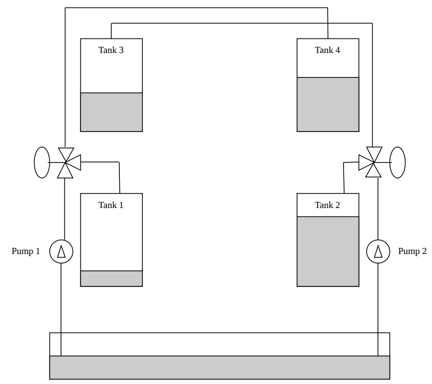

FIGURE 1 | **Quadruple-tank process.** Four interconnected tanks driven by two pumps; each pump i splits its flow by valve ratio γ_i between a lower tank and the diagonally opposite upper tank, producing the strong cross-coupling and, for $0 < \gamma_1 + \gamma_2 < 1$, the non-minimum-phase zero.

The quadruple-tank model follows these equations [27]:

$$\begin{aligned}
\frac{dh_1}{dt} &= -\frac{a_1}{A_1}\sqrt{2gh_1} + \frac{a_3}{A_1}\sqrt{2gh_3} + \frac{\gamma_1 k_1}{A_1}u_1 \\
\frac{dh_2}{dt} &= -\frac{a_2}{A_2}\sqrt{2gh_2} + \frac{a_4}{A_2}\sqrt{2gh_4} + \frac{\gamma_2 k_2}{A_2}u_2 \\
\frac{dh_3}{dt} &= -\frac{a_3}{A_3}\sqrt{2gh_3} + \frac{(1-\gamma_2)k_2}{A_3}u_2 \\
\frac{dh_4}{dt} &= -\frac{a_4}{A_4}\sqrt{2gh_4} + \frac{(1-\gamma_1)k_1}{A_4}u_1
\end{aligned} \tag{2}$$

In (2), tank and outlet cross-sectional areas $A = [28, 32, 28, 32]$, $a = [0.071, 0.057, 0.071, 0.057]$ are in cm^2; levels h_i (cm), pump voltages u_i (V), pump gains $k_1 = 3.14, k_2 = 3.29$ (cm^3/Vs), valve ratios are $\gamma_1 = 0.43, \gamma_2 = 0.34$, and $g = 981$ cm/s^2. The plant schematic is shown in Fig. 1.

3.3 | Control Objective

The plant configuration follows Johansson's NMP setup, described in [27]. The reference is not free to choose componentwise. With two pumps the plant has a two-dimensional equilibrium manifold: fixing the directly actuated pair (h_1, h_2) determines the steady inputs, and through them the upper levels (h_3, h_4), uniquely. Commanding $h_1 = h_2 = 10$ cm therefore forces

$$\bar{u} = \begin{pmatrix} \gamma_1 k_1 & (1-\gamma_2)k_2 \\ (1-\gamma_1)k_1 & \gamma_2 k_2 \end{pmatrix}^{-1} \begin{pmatrix} a_1\sqrt{2gh_1} \\ a_2\sqrt{2gh_2} \end{pmatrix}, \tag{3}$$

with $h_3 = \big((1-\gamma_2)k_2\bar{u}_2/a_3\big)^2/2g$ and $h_4 = \big((1-\gamma_1)k_1\bar{u}_1/a_4\big)^2/2g$, giving $\bar{u} \approx [2.61, 2.95]^T$ V and $h_3 = 4.16$, $h_4 = 3.44$ cm. The target state is accordingly set to $[10.0, 10.0, 4.16, 3.44]$, which lies on the manifold: no admissible input holds a reference off it, so an off-manifold target would make the upper-tank error metrics measure distance to an unreachable point rather than a control deficiency. Assumption 2 is thus satisfied by construction at nominal parameters, and the drift schedule below — which moves the manifold — is what makes the reference only approximately attainable thereafter. In addition to the challenge of NMP control, all controllers are subjected to structured adversarial perturbations: slow parametric drift of valve coefficients (modeling wear and aging), step-wise parameter deviations, and bounded external disturbances. These perturbations are non-vanishing, bounded and positioned outside of the nominal modeling assumptions typically required for LTV-MPC stability guarantees, while remaining physically plausible in industrial fluid systems. All perturbations are applied identically to all control systems, with no controller-specific tuning or compensation.

Performance is assessed over a 50-seed Monte-Carlo campaign (Sec. 5): each seed draws independent process- and measurement-noise realizations under an identical adversarial schedule,

- coefficient γ_1 of the first pump degrades at a constant rate of -0.0011/s;
- coefficient γ_2 of the second pump undergoes a step-wise 20% drop at $t = 25$ s.

A separate single-realization stress test (Sec. 5) additionally injects a 2.0 cm step disturbance into tank 2 at $t = 30$ s, on top of the drift schedule, to expose cross-coupled recovery under model mismatch.

4 | Controller Design & Stability Analysis

4.1 | Preliminaries: Kolmogorov–Arnold Representation

The architecture rests on the Kolmogorov–Arnold representation theorem [9, 10]: any continuous $g : [0,1]^n \to \mathbb{R}$ admits the finite superposition

$$g(x) = \sum_{q=0}^{2n} \Phi_q\Big(\sum_{p=1}^{n} \phi_{q,p}(x_p) \Big), \tag{4}$$

with continuous univariate inner and outer functions $\phi_{q,p}, \Phi_q$. A KAN [11] parameterizes these univariate maps as learnable splines placed on the network edges, while nodes perform summation; this is the structural dual of an MLP, in which fixed nonlinearities sit on nodes and learnable weights on edges. After offline training on input–output data of the plant nonlinearity, the trained univariate maps are distilled into a fixed, finite dictionary of basis functions, collected in the regressor $\Phi_{\mathrm{KAN}}(\cdot)$, so that the control-relevant nonlinearity of (1) is represented as a linear-in-parameters model,

$$f(x) = \Phi_{\mathrm{KAN}}(x)^T \theta^* + \varepsilon(x), \tag{5}$$

where θ^* are the (unknown, constant) ideal weights and ε is the residual approximation error. Only the linear weights θ are adapted online; the distilled basis is frozen. For the quadruple-tank instance, the inputs to Φ_{KAN} are the measured/estimated tank levels and the regressor reproduces the square-root outflow and cross-coupling terms of (2); the distilled law for u_1 reduces to the explicit algebraic form reported in Sec. 4.5.1.

Proposition 1 (Regularity of the distilled basis)**.** *The distilled regressor $\Phi_{\mathrm{KAN}} : \Omega \to \mathbb{R}^p$ is C^∞ and, on the compact operating set Ω, both bounded and Lipschitz: the constants $\Phi_{\max} := \sup_{x\in\Omega} \|\Phi_{\mathrm{KAN}}(x)\|$ and $L_\Phi := \sup_{x\in\Omega} \|\partial\Phi_{\mathrm{KAN}}/\partial x\|$ are finite, and $\|\Phi_{\mathrm{KAN}}(x) - \Phi_{\mathrm{KAN}}(y)\| \le L_\Phi \|x - y\|\ \forall\, x, y \in \Omega$. For the quadruple-tank instance $p = 10$: seven affine state/error features, two quadratic NMP-coupling features, and a bias term, as instantiated in (14)–(15).*

Proof. By the symbolic distillation (14)–(15), every component of Φ_{KAN} is either an affine function of the normalized features, the square of an affine function (the quadratic NMP-coupling terms), or a constant; hence each is a multivariate polynomial of degree at most two and $\Phi_{\mathrm{KAN}} \in C^\infty(\mathbb{R}^n, \mathbb{R}^p)$. Input normalization confines the operating envelope to the compact box Ω, on which the continuous map Φ_{KAN} attains its supremum (Weierstrass), so $\Phi_{\max} < \infty$. The Jacobian $\partial\Phi_{\mathrm{KAN}}/\partial x$ is affine, hence continuous and bounded on Ω; since the box Ω is convex, the mean-value inequality gives the stated Lipschitz bound with $L_\Phi < \infty$. Boundedness therefore needs no separate enforcement in implementation: it is a property of evaluating a fixed polynomial dictionary on a compact domain. □

Proposition 2 (Boundedness of ideal weights)**.** *Define ideal weights as $L^2(\Omega)$ projection of the plant nonlinearity onto the frozen basis,*

$$\theta^* := \arg\min_{\theta\in\mathbb{R}^p} \left\| f - \Phi_{\mathrm{KAN}}^T \theta \right\|^2_{L^2(\Omega)}, \tag{6}$$

which makes precise the ideal weights of (5) and pins $\varepsilon = f - \Phi_{\mathrm{KAN}}^T \theta^*$ as the $L^2(\Omega)$-orthogonal residual — the model error that no choice of weights can cancel. Let the distilled features be linearly independent on Ω, equivalently $\lambda_{\min}(G) > 0$ for the Gram matrix $G := \int_\Omega \Phi_{\mathrm{KAN}} \Phi_{\mathrm{KAN}}^T\, dx$. Then θ^* is unique and bounded:

$$\begin{aligned} \theta^* &= G^{-1} \int_\Omega \Phi_{\mathrm{KAN}}\, f\, dx, \\ \|\theta^*\| &\le \frac{\Phi_{\max}\ \sup_\Omega |f|\, |\Omega|}{\lambda_{\min}(G)} =: \theta_{\max}. \end{aligned} \tag{7}$$

For the vector plant the bound applies to each output channel, all sharing the Gram G.

Proof. The quadratic objective (6) has normal equations $G\theta^* = \int_\Omega \Phi_{\mathrm{KAN}} f\, dx$. Under $\lambda_{\min}(G) > 0$ the symmetric Gram G is positive definite, so the minimizer $\theta^* = G^{-1} \int_\Omega \Phi_{\mathrm{KAN}} f\, dx$ is unique and ε is $L^2(\Omega)$-orthogonal to the basis. Hence

$$\begin{aligned} \|\theta^*\| &\le \|G^{-1}\|_2 \left\| \int_\Omega \Phi_{\mathrm{KAN}} f\, dx \right\| \\ &\le \frac{1}{\lambda_{\min}(G)} \int_\Omega \|\Phi_{\mathrm{KAN}}\|\, |f|\, dx \le \frac{\Phi_{\max}\ \sup_\Omega |f|\, |\Omega|}{\lambda_{\min}(G)}, \end{aligned} \tag{8}$$

using $\|G^{-1}\|_2 = \lambda_{\min}(G)^{-1}$, boundedness of f (f continuous on the compact Ω, plant class (1)), and $\Phi_{\max} < \infty$ (Proposition 1). The hypothesis is a pure conditioning requirement: were $\|\theta^*\|$ unbounded the normal equations would force $\lambda_{\min}(G) = 0$, i.e. a nonzero v with $\int_\Omega |\Phi_{\mathrm{KAN}}^T v|^2 dx = 0$ and hence a linear dependence $\Phi_{\mathrm{KAN}}^T v \equiv 0$ among the features — contradicting their independence. □

Remark 1. Reading the weight bound off the distilled law. In practice the constant is read directly off the frozen distilled coefficients: $\|\theta^*\| \approx 13.7$ for the primary channel (23.8 for the secondary), so $\theta_{\max} \approx 23.8$ — mirroring how the integral bound $\bar{d}$ is read off logged data (Lemma 2). The bound is on the *ideal* weights θ^* of (6), which are constant by construction; it is not a bound on the running estimate $\hat{\theta}(t)$, whose excursion from θ^* is the weight error $\tilde{\theta}$ confined by the residual set (37) rather than by $\theta_{\max}$. The conditioning $\lambda_{\min}(G) > 0$ holds on the reference-varying distillation set of Algorithm 1, where the empirical Gram $G_N = \frac{1}{N}\sum_i \Phi_{\mathrm{KAN}}(\xi_i)\Phi_{\mathrm{KAN}}(\xi_i)^T$ is full rank. At a single regulation setpoint, by contrast, the error features collapse onto the level features ($x_5 = c_1 - x_1$, $x_7 = c_3 - x_3$, $x_8 = c_4 - x_4$) and G_N drops to rank seven, so PE (persistent excitation) is lost — precisely the regime in which the covariance bounding of the RLS gain (Assumption 3) and Corollary 2 operate. The same feature-conditioning $\lambda_{\min}$ thus wears two hats: bounded below on the distillation set it makes θ^* well-posed, while its degeneracy at the operating point is what renders the covariance bound of Assumption 3 load-bearing rather than decorative.

4.2 | Standing Conditions

The following conditions are used; each is referenced at the point of use.

Assumption 1 (Input map)**.** B has full column rank at the operating point, so $B^+B = I$.

Lemma 1 (Lumped disturbance bound)**.** *The lumped disturbance $\delta(t) \triangleq \varepsilon_{\mathrm{tv}}(x) + r_{\mathrm{ua}}(t) + B\,r_{\mathrm{slew}}(t) + r_{\mathrm{ukf}}(t)$, collecting the time-varying approximation error, the underactuation residual, the actuation-realization residual of Sec. 4.5.1, and the state-estimation residual, is uniformly bounded, $\delta \in \mathcal{L}_\infty$. The constant component of the approximation error ε is denoted d^*, with $\|d^*\| \le d_{\max}$.*

Proof. Each of the four components is bounded on the compact set Ω. (i) The approximation error $\varepsilon = f - \Phi_{\mathrm{KAN}}^T\theta^*$ is continuous on Ω — f is continuous on the compact Ω (plant class (1)) and Φ_{KAN} is bounded (Proposition 1) and θ^* bounded (Proposition 2) — hence bounded, so its time-varying part $\varepsilon_{\mathrm{tv}} = \varepsilon - d^*$ is bounded. (ii) The underactuation residual $r_{\mathrm{ua}} = (BB^+ - I)(\cdot)$ is a fixed projector applied to bounded arguments on Ω, hence bounded. (iii) The actuation residual $B\,r_{\mathrm{slew}}$ is bounded: the applied command (12) and the adaptation target (13) are confined to the pump range $[0, u_{\max}]$ by saturation, and the bounded-gain RLS driven by that saturated target keeps $\hat{\theta}$ — hence the design command (11) — bounded on the compact Ω. (iv) The estimation residual r_{ukf} is Lipschitz in the estimation error e_o, which is uniformly bounded by the regional estimator ISS of Lemma 3 below; that lemma rests only on Assumption 4 and is independent of δ, so citing it here introduces no circularity. Aggregating the four finite bounds gives $\delta \in \mathcal{L}_\infty$. □

Assumption 2 (Reference)**.** $x_{\mathrm{ref}}, \dot{x}_{\mathrm{ref}}$ are feasible and bounded, so that Ω is compact and forward-invariant under the closed loop.

Assumption 3 (Bounded gain matrix)**.** The RLS gain $\Gamma(t)$ is kept uniformly bounded and uniformly positive definite, $0 \prec \gamma_{\min} I \preceq \Gamma(t) \preceq \gamma_{\max} I$. This is enforced in implementation by projecting the updated covariance onto the band at every adaptation step (Sec. 4.5.2); exponential forgetting alone does *not* supply the upper bound, since in unexcited directions no information contraction opposes the $1/\lambda$ inflation of (20).

Assumption 4 (Estimator regularity)**.** The joint unscented Kalman filter (UKF) estimator state $x_{\mathrm{UKF}} = [\,h_{1\ldots4}^T\ \gamma_{1,2}^T\,]^T$ obeys the augmented dynamics $\dot{x}_{\mathrm{UKF}} = F(x_{\mathrm{UKF}}, u) + G_w w + G_\gamma \dot{\gamma}$, where F stacks the level dynamics (2) with the random-walk parameter model ($\dot{\gamma} \equiv 0$ nominally), $G_\gamma\dot{\gamma}$ injects the true valve drift as an exogenous signal, and the measurement is linear in the levels, $y = Cx_{\mathrm{UKF}} + \nu$ with $C = [\,I_4\ \ 0_{4\times2}\,]$: all four tank levels are sensed, while the valve ratios $\gamma_{1,2}$ are unmeasured, so the sensor block does not act on the γ-block. Let $F_x(t) := \int_0^1 \partial_{x_{\mathrm{UKF}}} F(\hat{x}_{\mathrm{UKF}} + s\,e_o, u)\,ds$ be the averaged state-Jacobian along the estimation error $e_o := x_{\mathrm{UKF}} - \hat{x}_{\mathrm{UKF}}$ and $\Theta(\tau, t)$ the state-transition matrix it generates. On the compact envelope $\Omega \times \mathcal{U}$ (states $\times$ admissible inputs):

(i) (Covariance sandwich) the filter covariance P_o is uniformly bounded and positive definite, $\underline{p}\,I \preceq P_o(t) \preceq \bar{p}\,I$ with $\underline{p}, \bar{p} > 0$ (empirically $\underline{p} \approx 4.7\times10^{-3}$, $\bar{p} \approx 9.6\times10^{-2}$ over the benchmark runs, i.e. $\mathrm{cond}(P_o) \lesssim 20$);

(ii) (Uniform observability) the pair $(F_x(t), C)$ is uniformly completely observable along every admissible trajectory: $\exists\, T, \alpha, \bar{\alpha} > 0$ with $\alpha I \preceq \int_t^{t+T} \Theta^T(\tau,t)\,C^T R_\nu^{-1} C\,\Theta(\tau,t)\,d\tau \preceq \bar{\alpha} I$, where $R_\nu \succ 0$ is the measurement covariance;

(iii) (Bounded slopes and curvature) $\|F_x(t)\| \le \bar{a}$ and $\sup_\Omega \|\partial^2 F\| \le \bar{H} < \infty$, both finite because F is smooth on the compact Ω (the only nonsmooth candidate, $\partial_h^2\sqrt{h}$, stays bounded since the tanks never drain to $h = 0$);

(iv) (Rate-bounded drift) $\|\dot{\gamma}(t)\| \le \bar{\nu}_\gamma$ for a.e. t, off a finite set of jump instants.

Condition (ii) is the structural crux: although $C = [\,I_4\ \ 0_{4x2}\,]$ reads all four levels, it cannot see γ directly, so the valve ratios are observable only through the multiplicative input channel $B(\gamma)u$. At the equilibrium operating point $\bar{h} = [10, 10, 4.16, 3.44]^\top$, $\bar{\gamma} = [0.43, 0.34]^\top$ with the steady holding input $\bar{u} \approx [2.61, 2.95]^\top$ V of (3), the discrete observability matrix of (F_x, C) has full column rank 6 (smallest singular value $\sigma_{\min} \approx 4.9\times10^{-3}$), but collapses to rank 4 at $u \equiv 0$, losing exactly the two γ-coordinates. Sensing all four levels rather than the (h_1, h_2) pair of Johansson's default keeps this recovery well-conditioned — γ_1 enters both h_1 and h_4, γ_2 both h_2 and h_3, so each ratio is carried by two measured channels. Hence (ii) holds precisely when the inputs persistently excite that channel.

4.3 | LTV-MPC Design

Linear Time-Varying MPC (LTV-MPC) re-solves at each control step ($dt = 0.003$ s), re-linearizing about the current state estimate $\hat{x}_k$, but discretizes its prediction model at the coarser step $\Delta t_p = 4$ s so that the $N = 30$ horizon spans $N\Delta t_p = 120$ s. This decoupling is deliberate, and the horizon is set by the plant rather than by convenience: the dominant tank time constants are 50–70 s and the non-minimum-phase transmission zero has a time constant of roughly 69 s, so the horizon must span the inverse response. A horizon discretized at the 3 ms control step ($N\,dt = 90$ ms) is myopic by orders of magnitude, and even a 3 s horizon remains short enough that the controller moves in the wrong direction on this plant, drifting steadily away from a reachable equilibrium [14]. The baseline is therefore given a horizon that comfortably exceeds the inverse-response time constant, so that the comparisons of Sec. 5 are not measuring an artificially handicapped MPC. Linearizing (2) about the operating point gives the affine model $\dot{x} = Ax + Bu + d_{\text{lin}}$, where d_{lin} is the linearization offset, the state Jacobian is

$$A = \begin{pmatrix} -\frac{a_1}{A_1}\sqrt{\frac{g}{2h_1}} & 0 & \frac{a_3}{A_1}\sqrt{\frac{g}{2h_3}} & 0 \\ 0 & -\frac{a_2}{A_2}\sqrt{\frac{g}{2h_2}} & 0 & \frac{a_4}{A_2}\sqrt{\frac{g}{2h_4}} \\ 0 & 0 & -\frac{a_3}{A_3}\sqrt{\frac{g}{2h_3}} & 0 \\ 0 & 0 & 0 & -\frac{a_4}{A_4}\sqrt{\frac{g}{2h_4}} \end{pmatrix}, \tag{9}$$

and the input matrix, following Johansson [27], is

$$B = \begin{pmatrix} \frac{\gamma_1 k_1}{A_1} & 0 \\ 0 & \frac{\gamma_2 k_2}{A_2} \\ 0 & \frac{(1-\gamma_2)k_2}{A_3} \\ \frac{(1-\gamma_1)k_1}{A_4} & 0 \end{pmatrix}, \tag{10}$$

which has full column rank at the NMP operating point (Assumption 1). Cost weighting follows standard state-input scaling guidelines [2, 3], with the QP solved by OSQP [4]: prediction horizon $N = 30$, $Q = \text{diag}([40, 40, 5, 5])$, and $R = \text{diag}([0.001, 0.001])$. Horizon, cost weights, and constraints are identical across all runs, with no preemptive tuning.

4.4 | Offline KAN Distillation

The regressor Φ_{KAN} of (5) is obtained offline by distilling the LTV-MPC policy of Sec. 4.3 into a fixed symbolic basis [14]. The procedure (Algorithm 1) has three stages. First, the LTV-MPC is rolled out across the operating envelope and reference set, producing a labeled dataset $\mathcal{D} = \{(\xi_i, u_i^{\text{MPC}})\}$ of normalized state/error features ξ and optimal commands. Second, a KAN with learnable B-spline activations on its edges is trained to imitate the policy, $\min_\Theta \sum_i \|\text{KAN}_\Theta(\xi_i) - u_i^{\text{MPC}}\|^2$. Third, each learned univariate spline is replaced by its best-fit symbolic primitive and the network is composed and simplified, yielding a finite basis $\Phi_{\text{KAN}}(\xi) \in \mathbb{R}^p$ ($p = 10$). Following Salkimbayev [14], the symbolic step is taken to supply the *support* only: the coefficients θ^* are then re-estimated on that fixed support by convex least squares against the teacher, rather than read off the symbolic composition directly. Two properties motivate this. The read-out is otherwise a dominant error source in its own right, and — because the law is linear in θ — the elementary negative-feedback requirement that raising a level must not raise the pump filling it, $\partial u_j/\partial h_j \le 0$, is a set of affine inequalities in θ and can be imposed inside the same convex program on a dense sample of Ω. On the present support the unconstrained read-out violates that requirement on 17.9% of the operating box, and enforcing it costs 0.11 percentage points of open-loop accuracy; the coefficients in (14)–(15) are the constrained solution. The basis is then frozen; only θ adapts online (Sec. 4.5.2). Distillation converts the iterative QP into a constant-time symbolic evaluation, at an imitation loss that is not negligible and that Salkimbayev [14] quantifies rather than assumes; the present contribution makes this frozen-basis controller adaptive and equips it with the stability certificate

Algorithm 1 Offline KAN distillation of the LTV-MPC policy

Require: operating set Ω, reference set $\mathcal{R}$, LTV-MPC policy π_{MPC}, KAN architecture
Ensure: frozen basis Φ_{KAN}, initial weights θ^*
$\mathcal{D} \leftarrow \emptyset$
for sampled $(x, x_{\text{ref}}) \in \Omega \times \mathcal{R}$ **do**
  $\xi \leftarrow$ NORMALIZE(x, x_{ref}); $u \leftarrow \pi_{\text{MPC}}(x, x_{\text{ref}})$
  $\mathcal{D} \leftarrow \mathcal{D} \cup \{(\xi, u)\}$
end for
train KAN B-spline edges Θ: $\min_\Theta \sum_{\mathcal{D}} \|\text{KAN}_\Theta(\xi) - u\|^2$
$\Phi_{\text{KAN}} \leftarrow$ SYMBOLICEXTRACT(Θ) ▷ spline→primitive; compose; simplify: support only
$\theta^* \leftarrow \arg\min_\theta \sum_{\mathcal{D}} \|\Phi_{\text{KAN}}(\xi)^T\theta - u/u_{\max}\|^2$ s.t. $\partial u_j/\partial h_j \le 0$ on Ω ▷ convex re-estimation
freeze Φ_{KAN} ▷ only θ adapts online
return Φ_{KAN}, θ^*

of Sec. 4.6. Online adaptation of θ is in this light not only a response to drift but also the mechanism by which a fixed read-out error is absorbed after deployment, since any static mismatch between $\Phi_{\text{KAN}}^T\theta^*$ and the teacher policy lies in the span of the basis and is therefore correctable by the weights. The symbolic law is extracted on normalized inputs and rescaled to physical units at deployment, which ensures benchmark fairness and underlies the regressor bound $\Phi_{\max}$ of Proposition 1.

4.5 | Controller Design

4.5.1 | Structure of LCA-RLS-KAN

The controller is specified in a design form, in which the error dynamics and the Lyapunov analysis of Sec. 4.6 are conducted, and an implemented form, which the online loop actually evaluates. The design form combines the feedforward/adaptive law with the proportional-integral (PI) law,

$$u_{\text{raw}}(t) = B^+\Big(\dot{x}_{\text{ref}} - \Phi_{\text{KAN}}^T\hat{\theta} - K\,e(t) - \hat{d}(t)\Big), \tag{11}$$

where B^+ is the left pseudoinverse of B (Assumption 1; $B \in \mathbb{R}^{4\times 2}$ has full column rank at the NMP operating point), $\dot{x}_{\text{ref}}$ is the feedforward term, $\Phi_{\text{KAN}}^T\hat{\theta}$ cancels the learned nonlinearity (5), $K = K_P + \bar{K}_L$ is the total proportional gain (proportional gain K_P plus the Lipschitz-shaping gain $\bar{K}_L := (K_L/e_{\max})I$, see Sec. 4.5.2), and $\hat{d}(t)$ is the integral state with dynamics (21). The sign of the regressor and integral terms in (11) is the cancellation sign required by the error dynamics (27).

The implemented form evaluates (11) composed. Every static term in (11) is affine in the distilled feature dictionary of Sec. 4.4 — B^+ is a fixed matrix, $\dot{x}_{\text{ref}}$ is constant for the regulation task, and Ke is affine in the level and error features — so the static part of the composed law lies in the span of Φ_{KAN}, and the command applied to the plant is a single saturated dot product per pump,

$$u(t_k) = \text{sat}_{[0,u_{\max}]}\big(u_{\max}[\,\hat{\theta}_1^T\varphi_k,\ \hat{\theta}_2^T\varphi_k\,]^T\big), \tag{12}$$

where $\varphi_k = \Phi_{\text{KAN}}(\xi_k)$ is the regressor evaluated on the normalized features, $u_{\max} = 12$ V is the command denormalization scale, and the saturation confines the command to the physical pump range. The distilled polynomials (14)–(15) are exactly this composition at the initial weights; the composed coefficients are an affine function of the design weights $\hat{\theta}$ (with the integral state entering the bias coordinate), so the two forms carry the same adaptive degrees of freedom. The slowly varying terms of (11) — the integral state $\hat{d}$ and the corrective proportional action — are therefore not added on the actuator line at each instant; they enter through the adaptation channel, carried by the time-varying weights $\hat{\theta}(t)$. At every step a PI-corrected target is formed around the applied command — proportional action on each pump's estimated level error ($K_{\text{corr}} = 3.0$ V/cm, the implementation counterpart of the design gain K), the low-pass-filtered cross-coupled upper-tank error realizing $\hat{d}$ (filter constant $\alpha = 0.05$, weight 1.0 — i.e. the coupled-tank error enters each pump's correction with the same weight as that pump's own error, the untuned choice; the value is stated here because a smaller weight is appropriate only when the upper tanks carry an unattainable target, which (3) removes), and a saturating cross-coupling brake — and the target, not the actuator command, is rate-limited:

$$\begin{aligned} u_{\text{tgt}}(t_k) = \text{SlewLim}_{\Delta}\big(\tilde{u}_{\text{tgt}}(t_k)\big) &:= u_{\text{tgt}}(t_{k-1}) \\ &+ \text{sat}_{[-\Delta,\Delta]}\big(\tilde{u}_{\text{tgt}}(t_k) - u_{\text{tgt}}(t_{k-1})\big), \end{aligned} \tag{13}$$

with $\tilde{u}_{\text{tgt}}$ the unlimited PI-corrected target: the target may change by at most Δ volts per step ($\Delta = 0.1$ V here) and is saturated to the same $[0, u_{\max}]$ range. The forgetting-RLS of Sec. 4.5.2 regresses the weights onto the slew-limited target every T_a steps, taking as its measurement the target in the normalized units of the regressor, $u_{\text{tgt}}/u_{\max}$, so that measurement and prediction $\Phi^T\hat{\theta}$ carry the same scale as required by the model $y = \Phi^T\theta + \epsilon$ of (19); the innovation is then the normalized command error $(u_{\text{tgt}} - u_{\text{raw}})/u_{\max}$. The applied command therefore tracks the slew-limited target through the low-gain adaptation rather than through a hard constraint on the actuator line, bounding actuator slew and fatigue while keeping the actuation path branch-free. The realization residual $r_{\text{slew}}(t) := u(t) - u_{\text{raw}}(t)$ between the implemented command (12) and the design command (11) is uniformly bounded — u and u_{tgt} are confined to $[0, u_{\max}]$ by construction, and the bounded-gain RLS driven by this saturated target keeps $\hat{\theta}$, hence u_{raw}, bounded on the compact Ω [23, 18] — and enters the loop additively through $B\, r_{\text{slew}}$, absorbed into the lumped disturbance δ (Lemma 1). The previously monolithic u_L is thus resolved into two well-posed objects: a proportional contribution $-\bar{K}_L e$ folded into K, and the bounded rate saturation (13); this removes the dimensional inconsistency of treating a slew limit as a length.

To demonstrate the transparency of the KAN architecture, the distilled control law for u_1 and u_2 extracts into a single algebraic form (in normalized units; the deployed command is rescaled by $u_{\max}$ and saturated per (12)). What is displayed below is the law *at initialization*, $\hat{\theta}(0) = \theta^*$: the basis is frozen for all time, while the ten linear coefficients per pump are exactly the quantities the RLS adapts online (Sec. 4.5.2), and they move several-fold over a run (Sec. 5.4). The transparency claim is therefore a structural one — the functional form is fixed, known in closed form, and never changes, so the deployed controller is at every instant an explicit degree-two polynomial whose current coefficients can be read off and audited — rather than a claim that the particular numbers below remain the operative gains:

$$\begin{aligned} u_{1,\text{raw}} &= -3.061817\, x_1 - 0.914763\, x_2 - 6.955925\, x_3 \\ &- 3.407967\, x_4 + 6.469391\, x_5 - 8.475651\, x_7 \\ &- 1.330975\, x_8 + 0.025392\, (-10.000158\, x_1 \\ &- 9.999930)^2 + 0.073162\, (-1.670916\, x_1 \\ &+ 3.952262\, x_3 - 3.957915\, x_4 + 0.059323\, x_8 \\ &- 1.277487)^2 - 1.747098 \end{aligned} \tag{14}$$

$$\begin{aligned} u_{2,\text{raw}} &= -5.403939\, x_1 - 1.500293\, x_2 - 13.710084\, x_3 \\ &- 5.264385\, x_4 + 11.257138\, x_5 - 11.914867\, x_7 \\ &- 6.449098\, x_8 + 0.047566\, (-10.000158\, x_1 \\ &- 9.999930)^2 - 0.022053\, (-1.670916\, x_1 + \\ &3.952262\, x_3 - 3.957915\, x_4 + 0.059323\, x_8 \\ &- 1.277487)^2 - 3.109220 \end{aligned} \tag{15}$$

where x_1 is the normalized height of tank 1 (h_1), x_2 is the normalized height of tank 2 (h_2), x_3 is the normalized height of tank 3 (h_3), x_4 is the normalized height of tank 4 (h_4), x_5 is the normalized error of tank 1 ($h_{1target} - h_1$), x_6 is the normalized error of tank 2 ($h_{2target} - h_2$), x_7 is the normalized error of tank 3 ($h_{3target} - h_3$), x_8 is the normalized error of tank 4 ($h_{4target} - h_4$).

The coefficients in (14)–(15) reveal features structurally unavailable in optimization-based controllers: quadratic terms carry the NMP coupling in the respective pumps and the negative linear terms carry physical damping adjacent to setpoint deviation. The two modes hence provide a multilayered control law, which is specifically designed to handle the inverse response characteristic of the NMP plant configuration. This structural reading persists under adaptation, since the terms it refers to are basis elements rather than particular coefficient values; what adaptation changes is the weight each term carries.

The modifications ensuring boundedness and ISS are described next.

4.5.2 | Lipschitz-Continuous RLS Update Law

On top of the baseline KAN-distilled control law, three modifications enable safe online adaptation.

1) Weight adaptation. The linear weight vector $\hat{\theta}$ is updated by the continuous-time RLS law with forgetting factor λ [18, 24]:

$$\dot{\hat{\theta}} = \Gamma\, \Phi\, P\, e, \tag{16}$$

$$\dot{\Gamma} = \rho\, \Gamma - \Gamma\, \Phi\, \Phi^T \Gamma, \qquad \rho \approx \frac{1-\lambda}{T_a}, \tag{17}$$

where $\Phi \equiv \Phi_{\text{KAN}}(x)$ abbreviates the frozen regressor, T_a is the adaptation interval, and $0 < \lambda < 1$. The distilled nonlinear basis is frozen; only the linear weights adapt, so $\tilde{\theta} = \hat{\theta} - \theta^*$ is the weight estimation error (with θ^* constant, $\dot{\tilde{\theta}} = \dot{\hat{\theta}}$) [23]. The discrete realization used in the implementation is the Joseph-stabilized forgetting-RLS recursion; for the per-pump target $y_k = \Phi_k^T\theta + \epsilon_k$,

$$L_k = \Gamma_{k-1}\Phi_k\big(\lambda R_k + \Phi_k^T\Gamma_{k-1}\Phi_k\big)^{-1}, \tag{18}$$

$$\hat{\theta}_k = \hat{\theta}_{k-1} + L_k\big(y_k - \hat{\theta}_{k-1}^T\Phi_k\big), \tag{19}$$

$$\Gamma_k = \tfrac{1}{\lambda}\big[(I - L_k\Phi_k^T)\Gamma_{k-1}(I - L_k\Phi_k^T)^T + L_k R_k L_k^T\big], \tag{20}$$

with R_k being the measurement-noise covariance; the discrete covariance Γ_k is the sampled counterpart of the continuous gain Γ in (17) (with $\rho \approx \frac{1-\lambda}{T_a}$), and the Joseph form (20) keeps it symmetric positive definite. To prevent covariance wind-up when PE is absent, Γ is kept within $[\gamma_{\min} I, \gamma_{\max} I]$ (Assumption 3) by projecting the eigenvalues of Γ_k onto that band after each update; constant forgetting alone does not guarantee an upper bound on Γ without excitation. The necessity is quantitative: on the single-setpoint benchmark the regressor leaves three of the ten feature directions unexcited (Sec. 5), and in those directions the recursion (20) inflates the covariance by exactly $1/\lambda$ per update, so the largest eigenvalue of Γ grows geometrically — reaching 5.30×10^3 from an initial 10^3 over the 50 s horizon, i.e. precisely the pure-forgetting factor λ^{-N_a} with $N_a = 1667$ adaptation steps, confirming that no contraction acts there. Without the projection the band would be enforced only by the finite horizon. The bounds used are $\gamma_{\max} = 10^4$ and $\gamma_{\min} = 10^{-9}$, both set outside the measured envelope (the observed extremes are 5.30×10^3 and 1.19×10^{-8}), so the projection is inactive on the reported runs and the results of Sec. 5 are unaffected; its role is to make the band a property of the update law rather than of the run length. The upper bound is the load-bearing one, entering the damping condition (33) through $\lambda_{\min}(\Gamma^{-1}) = \gamma_{\max}^{-1}$. The lower bound is deliberately loose: the covariance genuinely contracts towards 10^{-8} in the excited directions, so $\gamma_{\min}$ here guarantees invertibility of Γ rather than a tight conditioning number, and the resulting spread makes the certificate of Sec. 4.6 a qualitative boundedness statement rather than a numerically tight bound.

2) Bounded integral action. The integral state $\hat{d}$ regulates the cumulative deviation from the *setpoint*; its design law is

$$\dot{\hat{d}} = K_I\, e, \qquad K_I = k_I I \succ 0, \tag{21}$$

with the band $[-\bar{d}, \bar{d}]$ kept forward-invariant for $\hat{d}$ (Lemma 2). In implementation this band is not enforced by an active conditional-integration clamp: $\hat{d}$ is realized as the low-pass-filtered error of Sec. 4.5.1 (filter constant α), a leaky accumulation whose output is a convex combination of past errors and is therefore confined to the error range of the compact envelope by construction. The constraint — integration frozen whenever it would push $\hat{d}$ outside $[-\bar{d}, \bar{d}]$ — is the corresponding guardrail for pure-integrator realizations of (21) and is subsumed by the filter structure here. Unlike the weight covariance Γ, the integral gain K_I is constant, so the design update (21) carries no analogue of the forgetting-induced metric decay that damps $\tilde{\theta}$ (Sec. 4.6); boundedness of $\hat{d}$, and hence of $\tilde{d} = \hat{d} - d^*$, is therefore supplied directly by the construction of the update — filter leak or constraint — rather than by the Lyapunov argument. This also removes any need to tie K_I to P: any SPD K_I may be chosen freely for transient shaping.

3) Lipschitz shaping and slew limiting. Lipschitz continuity of the command is enforced through the deviation-dependent proportional gain $\bar{K}_L = (K_L/e_{\max})I$ (folded into K in (11)) together with the per-step rate saturation (13) of the adaptation target [26]; the applied command (12) is itself Lipschitz in the state, being a saturated polynomial evaluation on the compact Ω (Proposition 1). This keeps the control law restrained and tractable while bounding actuator fatigue.

Lemma 2 (Boundedness of the integral state). *Let the integral update (21) be realized so that the band $[-\bar{d}, \bar{d}]$ is forward-invariant for $\hat{d}$ — by the leaky low-pass realization implemented here, whose output is a convex combination of past errors, or by a conditional-integration (anti-windup) clamp. Then $\|\hat{d}(t)\| \le \bar{d}$ for an explicit, known $\bar{d}$, $\forall t$, by construction of the update rather than as a consequence of the Lyapunov argument of Sec. 4.6. Together with $\|d^*\| \le d_{\max}$ (Lemma 1), the integral-state error obeys the uniform bound $\|\tilde{d}(t)\| \le \bar{d} + d_{\max} =: \bar{d}_{\mathrm{tot}}\ \forall t$.*

Proof. Under the implemented leaky realization the discrete update is the convex combination $\hat{d}^+ = (1-\alpha)\hat{d} + \alpha\, e$: with $\hat{d}(0) = 0$, coordinatewise $|\hat{d}_i(t)| \le \max_{\tau \le t} |e_i(\tau)|$, and since the error is confined to the compact envelope Ω (Assumption 2), any $\bar{d} \ge \sup_\Omega \|e\|_\infty$ renders the band positively invariant a priori. Under a clamp, integration acts coordinatewise on the accumulation (21): whenever an increment $K_I e\, \Delta t$ would carry $\hat{d}$ outside $[-\bar{d}, \bar{d}]$, it is suppressed and $\hat{d}$ is held at the boundary. In either case $\|\hat{d}(t)\| \le \bar{d}$ for all t, independently of any Lyapunov estimate. With $\tilde{d} = \hat{d} - d^*$ and $\|d^*\| \le d_{\max}$ (Lemma 1), the triangle inequality gives $\|\tilde{d}\| \le \bar{d} + d_{\max} =: \bar{d}_{\mathrm{tot}}$. Empirically the tighter band $\bar{d} = 2.5$ cm is certified by the logs: across the Monte-Carlo campaign of Sec. 5 the integral coordinate never exceeds 2.43 cm. □

The complete online cycle — estimation, constant-time inference, saturated actuation with rate-limited target shaping, and periodic weight adaptation — is summarized in Algorithm 2.

Algorithm 2 Online LCA-RLS-KAN control step at time t_k

Require: frozen Φ_{KAN}, scale $u_{\max}$, weights $\hat{\theta}_1, \hat{\theta}_2$, covariances Γ_1, Γ_2, forgetting λ, slew Δ, interval T_a, UKF

UKF predict with $u(t_{k-1})$; update with $z_k \Rightarrow \hat{x}_k = [\hat{h}, \hat{\gamma}]$

$\varphi \leftarrow \Phi_{\mathrm{KAN}}(\hat{x}_k, x_{\mathrm{ref}})$ ▷ $\mathcal{O}(p)$, branch-free

$u_{\mathrm{raw}} \leftarrow u_{\max}[\hat{\theta}_1^\top \varphi,\ \hat{\theta}_2^\top \varphi]$ ▷ constant-time inference

$u(t_k) \leftarrow \mathrm{sat}_{[0,u_{\max}]}(u_{\mathrm{raw}}(t_k))$ (12); apply to plant

form PI-corrected target $\tilde{u}_{\mathrm{tgt}}$ (filtered error + cross-coupling brake); $u_{\mathrm{tgt}} \leftarrow \mathrm{SlewLim}_\Delta(\tilde{u}_{\mathrm{tgt}})$ (13)

if $k \bmod T_a = 0$ then

  for $j \in \{1, 2\}$ do

    $(\hat{\theta}_j, \Gamma_j) \leftarrow$ FF-RLS$(\hat{\theta}_j, \Gamma_j, \varphi, u_{\mathrm{tgt},j}/u_{\max}; \lambda)$ ▷ target in normalized units

  end for

end if

4.5.3 | Lyapunov Candidate

The energy function is taken over the tracking and weight errors only,

$$V(\zeta) = \tfrac{1}{2} e^T P e + \tfrac{1}{2} \tilde{\theta}^T \Gamma^{-1} \tilde{\theta}, \tag{22}$$

with $\zeta = [\, e^T\ \tilde{\theta}^T\,]^T$, $e = x - x_{\mathrm{ref}}$ and $\tilde{\theta}$ the weight estimation error. The integral-state error $\tilde{d} = \hat{d} - d^*$ is deliberately excluded from V: by Lemma 2 it is already uniformly bounded by construction of its band-confined realization, independently of any Lyapunov argument, and it is instead carried into the analysis below as a bounded perturbation acting on the $(e, \tilde{\theta})$-subsystem. Since P and Γ^{-1} (Assumption 3) are symmetric positive definite, V is positive definite and radially unbounded, and there exist $c_1, c_2 > 0$ with

$$c_1\|\zeta\|^2 \le V(\zeta) \le c_2\|\zeta\|^2, \tag{23}$$

i.e. V satisfies class-$\mathcal{K}_\infty$ bounds on the same vector ζ that the dissipation rate of Sec. 4.6 will control (cf. Khalil and Grizzle [26, Lemma 4.3]).

4.6 | Stability Analysis

Differentiating (22) along the closed-loop trajectories,

$$\dot V = e^T P\dot e + \tilde\theta^T\Gamma^{-1}\dot{\tilde\theta} + \tfrac{1}{2}\tilde\theta^T \tfrac{d}{dt}(\Gamma^{-1})\tilde\theta. \tag{24}$$

With the plant model (5), split the approximation error into its constant and time-varying parts, $\varepsilon = d^* + \varepsilon_{\mathrm{tv}}(x)$, with d^* being the constant bias. Substituting the command (11)–(13) into $\dot e$, we get:

$$\dot e = \dot x - \dot x_{\mathrm{ref}} = f(x) + Bu - \dot x_{\mathrm{ref}}; \tag{25}$$

using $BB^+ \approx I$, collecting the time-varying terms into the lumped disturbance δ:

$$\delta = \varepsilon_{\mathrm{tv}} + r_{\mathrm{ua}} + B\, r_{\mathrm{slew}} + r_{\mathrm{ukf}} \tag{26}$$

of Lemma 1 (so the approximation-error remainder $\varepsilon - d^* = \varepsilon_{\mathrm{tv}}$ is only one of its four components), and setting $\theta^* - \hat\theta = -\tilde\theta$, $d^* - \hat d = -\tilde d$, the error dynamics are

$$\dot e = -\Phi^T\tilde\theta - \tilde d - Ke + \delta(x), \qquad K = K_P + \bar K_L. \tag{27}$$

Group the last two exogenous terms into a single effective disturbance:

$$\omega(t) := \delta(x(t), t) - \tilde d(t), \tag{28}$$

so that $\dot e = -\Phi^T\tilde\theta - Ke + \omega$. Because $\tilde d$ is bounded by construction (Lemma 2) rather than by anything established below, ω may be treated as an exogenous bounded signal acting on the $(e, \tilde\theta)$-subsystem: $\|\omega(t)\| \le \|\delta\|_\infty + \bar d_{\mathrm{tot}} =: \bar\omega < \infty$, uniformly in t. Since θ^* is constant, $\dot{\tilde\theta} = \dot{\hat\theta}$ is given by (16). The gain law (17) yields the inverse-covariance derivative

$$\tfrac{d}{dt}(\Gamma^{-1}) = \Phi\Phi^T - \rho\,\Gamma^{-1}. \tag{29}$$

Substituting (27)–(29) into (24), the weight cross-term $\tilde\theta^T\Gamma^{-1}(\Gamma\Phi P e) = \tilde\theta^T\Phi P e$ cancels $-e^T P\Phi^T\tilde\theta$ exactly (both equal $e^T P\Phi^T\tilde\theta$ by symmetry of P), leaving

$$\dot V = -e^T PKe + e^T P\omega + \tfrac{1}{2}\tilde\theta^T\Phi\Phi^T\tilde\theta - \tfrac{1}{2}\rho\,\tilde\theta^T\Gamma^{-1}\tilde\theta. \tag{30}$$

No cross-term involving $\tilde d$ remains to be cancelled or bounded here, because $\tilde d$ was never placed in V in the first place (Sec. 4.5.3); its effect on the loop enters only through the bounded forcing term ω in (30), exactly as δ does. The cross/indefinite terms are bounded by Cauchy–Schwarz and Young's inequalities. For the disturbance term, with $\mu > 0$,

$$e^T P\omega \le \lambda_{\max}(P)\,\|e\|\,\|\omega\| \le \tfrac{\mu}{2}\|e\|^2 + \tfrac{\lambda_{\max}(P)^2}{2\mu}\|\omega\|^2. \tag{31}$$

By Proposition 1, $\frac{1}{2}\tilde\theta^T\Phi\Phi^T\tilde\theta \le \frac{1}{2}\Phi_{\max}^2\|\tilde\theta\|^2$, and combining with the forgetting term,

$$\begin{aligned}&\tfrac{1}{2}\Phi_{\max}^2\|\tilde\theta\|^2 - \tfrac{1}{2}\rho\,\tilde\theta^T\Gamma^{-1}\tilde\theta \le \\ &-\tfrac{1}{2}(\rho\,\lambda_{\min}(\Gamma^{-1}) - \Phi_{\max}^2)\|\tilde\theta\|^2.\end{aligned} \tag{32}$$

This term is exactly the "implicit leakage" of exponential forgetting: Γ^{-1} decaying at rate ρ contributes a genuine negative-definite term to $\dot V$ through the metric itself (not through any explicit pull-to-zero in $\dot{\hat\theta}$), provided Γ^{-1} stays bounded away from 0 (Assumption 3). The design integrator (21) has no such mechanism — K_I is constant, so no analogous term arises for $\tilde d$; this is precisely why $\tilde d$ is handled by construction (Lemma 2) rather than by the Lyapunov dissipation; a leak placed on $\hat d$ itself is what the companion σ-modified design of Sec. 4.7 would convert into dissipation.

Taking the gains scalar, $K = kI$ with $k = k_P + \bar k_L > 0$, so $PK = kP \succ 0$ and $\lambda_{\min}(PK) = k\,\lambda_{\min}(P)$ (more generally, it suffices that $PK + K^T P \succ 0$). Requiring the error term to dominate the disturbance cross-term, $\lambda_{\min}(PK) > \mu/2$; choose $\mu \in (0, 2\lambda_{\min}(PK))$, set $\mu = \lambda_{\min}(PK)$. By Assumption 3, Γ^{-1} is bounded with $\lambda_{\min}(\Gamma^{-1}) > 0$, and the damping condition

$$\rho\,\lambda_{\min}(\Gamma^{-1}) > \Phi_{\max}^2 \tag{33}$$

is enforceable since $\Phi_{\max}$ is finite (Proposition 1). Collecting (31)–(32), the time derivative satisfies

$$\begin{aligned}\dot V \le& -\Big(\lambda_{\min}(PK) - \tfrac{\mu}{2}\Big)\|e\|^2 \\ &-\tfrac{1}{2}\Big(\rho\,\lambda_{\min}(\Gamma^{-1}) - \Phi_{\max}^2\Big)\|\tilde\theta\|^2 + \tfrac{\lambda_{\max}(P)^2}{2\mu}\|\omega\|^2.\end{aligned} \tag{34}$$

This is a negative-definite term in every component of $\zeta = (e, \tilde\theta)$, with a forcing term bounded uniformly by $\bar\omega$ (Lemma 2); the integral channel contributes only through this bounded forcing, never through an unbounded or unaccounted state.

Define

$$\begin{aligned}\alpha &:= \min\Big\{\lambda_{\min}(PK) - \tfrac{\mu}{2},\ \tfrac{1}{2}\big(\rho\lambda_{\min}(\Gamma^{-1}) - \Phi_{\max}^2\big)\Big\}, \\ \gamma &:= \tfrac{\lambda_{\max}(P)^2}{2\mu},\end{aligned} \tag{35}$$

both positive under (33). Then (34) reads $\dot V \le -\alpha\|\zeta\|^2 + \gamma\|\omega\|^2$, a dissipation inequality on the same state $\zeta = [\,e^T\ \tilde\theta^T\,]^T$ bounded in (23), with ω as in (28) satisfying $\|\omega\|_\infty \le \bar\omega$. The UKF estimation residual is treated as a component of δ (Lemma 1) and handled by an interconnection (cascade-ISS) argument rather than a separation principle: the estimator subsystem is shown regionally ISS in Lemma 3, and its interconnection with the controller subsystem is closed in Lemma 4.

Theorem 1 (Stability of LCA-RLS-KAN).
Consider the closed-loop system formed by the plant (1) and the LCA-RLS-KAN controller (11)–(13), (16)–(21), under Proposition 1, Assumptions 1–3, and Lemmas 1 and 2. Let the gains P, Γ be SPD, $K = kI \succ 0$, and let $\hat d$ be realized so that the band invariance of Lemma 2 holds (leaky low-pass realization or anti-windup clamp). If the damping condition (33) holds, then with α, γ in (35) and ω as in (28),

$$\dot V \le -\alpha\,\|\zeta\|^2 + \gamma\,\|\omega\|^2. \tag{36}$$

By (23), V is an ISS-Lyapunov function in the sense of Sontag [25] (cf. Khalil and Grizzle [26, Def. 4.7, Thm. 4.19]); hence the $(e,\tilde{\theta})$-subsystem is input-to-state stable with respect to ω, and $\zeta = (e,\tilde{\theta})$ is uniformly ultimately bounded, converging to the residual set

$$\mathcal{R} = \left\{\zeta \,:\, \|\zeta\| \le \sqrt{\tfrac{c_2}{c_1}\tfrac{\gamma}{\alpha}}\,\bar{\omega}\right\}. \tag{37}$$

The integral-state error $\tilde{d}$ is bounded independently by $\bar{d}_{\text{tot}}$ (Lemma 2), so the full triple $(e,\tilde{\theta},\hat{d})$ remains bounded $\forall t$.

Proof. The dissipation inequality (36) was assembled in (24)–(35); it remains to convert it into an explicit bound. By the lower bound in (23), $\|\zeta\|^2 \ge V/c_2$, so (36) with $\|\omega\|_\infty \le \bar{\omega}$ gives the scalar comparison inequality

$$\dot{V} \le -\eta V + \gamma\bar{\omega}^2, \qquad \eta := \frac{\alpha}{c_2} > 0. \tag{38}$$

The comparison lemma [26, Lemma 3.4] integrates (38) to

$$V(t) \le e^{-\eta t}V(0) + \frac{\gamma}{\eta}\left(1 - e^{-\eta t}\right)\bar{\omega}^2 \le e^{-\eta t}V(0) + \frac{c_2\gamma}{\alpha}\bar{\omega}^2. \tag{39}$$

Applying $c_1\|\zeta\|^2 \le V(t)$ and $V(0) \le c_2\|\zeta(0)\|^2$ from (23) together with $\sqrt{a+b} \le \sqrt{a} + \sqrt{b}$,

$$\|\zeta(t)\| \le \underbrace{\sqrt{\tfrac{c_2}{c_1}}\,\|\zeta(0)\|\,e^{-\eta t/2}}_{\beta(\|\zeta(0)\|,\,t)\in\mathcal{KL}} + \underbrace{\sqrt{\tfrac{c_2}{c_1}\tfrac{\gamma}{\alpha}}\,\bar{\omega}}_{\kappa(\|\omega\|_\infty)\in\mathcal{K}}. \tag{40}$$

Estimate (40) is exactly an ISS bound — a class-$\mathcal{KL}$ transient in $(\|\zeta(0)\|, t)$ plus a class-$\mathcal{K}$ function of $\|\omega\|_\infty$ [25], [26, Thm. 4.19]. Hence the $(e,\tilde{\theta})$-subsystem is regionally ISS with respect to ω, ζ is UUB, and

$$\limsup_{t\to\infty} \|\zeta(t)\| \le \sqrt{\tfrac{c_2}{c_1}\tfrac{\gamma}{\alpha}}\,\bar{\omega}, \tag{41}$$

i.e. ζ enters the residual set (37), approached at exponential rate $\eta/2 = \alpha/(2c_2)$. The integral coordinate is bounded separately, $\|\tilde{d}(t)\| \le \bar{d}_{\text{tot}}$ for all t by Lemma 2, so $(e,\tilde{\theta},\hat{d})$ is bounded while $(e,\tilde{\theta})$ is UUB. □

Remark 2. Scope of the certificate. The conclusion of Theorem 1 is regional rather than global, and the distinction is worth stating precisely, because for this plant the region is not a restriction. The constants the certificate rests on — $\Phi_{\max}$ and L_Φ (Proposition 1), $\theta_{\max}$ (Proposition 2), and $\|\delta\|_\infty$ (Lemma 1) — are suprema over the compact operating envelope Ω, and a global statement is unavailable for a structural reason, not a technical one: the distilled features are polynomials of degree two (Proposition 1), so $\Phi_{\max}$ is finite only on a bounded set, and the damping condition (33) — which is what renders $\alpha > 0$ — cannot be enforced on all of $\mathbb{R}^n$ by any choice of gains. A globally valid certificate would require a globally bounded regressor, which the symbolic read-out that supplies the interpretability claim does not provide.

What replaces a global claim here is that Ω can be taken to be the entire physically realizable state set. The tank levels are confined to $[0, h_{\max}]$ by finite vessel capacity and the commands to $[0, u_{\max}]$ by pump saturation (12), so $\Omega \times \mathcal{U}$ is forward-invariant by construction of the plant rather than by hypothesis on the closed loop (Assumption 2, Lemma 4); states outside it are not merely uncertified but unrealizable — indeed the vector field (2) is not defined for $h < 0$. The domain of attraction certified by (40) therefore contains every initial condition the plant can physically occupy, and the residual set (37) is a subset of Ω. The estimate is regional in form and exhaustive in scope: no admissible initial condition is excluded.

Lemma 3 (Regional ISS of the joint estimator)**.** *Under Assumption 4, realize the UKF as the output-injection observer*

$$\begin{aligned} \dot{\hat{x}}_{\text{UKF}} &= F(\hat{x}_{\text{UKF}}, u) + K_{\text{f}}(t)\left(y - C\hat{x}_{\text{UKF}}\right), \\ K_{\text{f}} &= P_o C^T R_\nu^{-1}. \end{aligned} \tag{42}$$

Then the estimation error $e_o = x_{\text{UKF}} - \hat{x}_{\text{UKF}}$ is regionally ISS with respect to $(\nu, \dot{\gamma})$, with gains uniform over $\Omega \times \mathcal{U}$:

$$\|e_o(t)\| \le \beta_o(\|e_o(0)\|, t) + \rho_o(\|(\nu,\dot{\gamma})\|_\infty) + b_o, \tag{43}$$

for some $\beta_o \in \mathcal{KL}$, $\rho_o \in \mathcal{K}$, and a bias floor $b_o = \mathcal{O}(\bar{H}\bar{p})$ induced by unscented statistical linearization (vanishing as the filter is tightened, $\bar{p}\downarrow$, or the plant curvature is mild, $\bar{H}\downarrow$).

Proof. Subtracting (42) from the augmented plant and using the mean-value form $F(x_{\text{UKF}}, u) - F(\hat{x}_{\text{UKF}}, u) = F_x(t)\,e_o$ (exact, with F_x as in Assumption 4) gives

$$\dot{e}_o = \left(F_x - K_{\text{f}}C\right)e_o + \Delta F_x\, e_o - K_{\text{f}}\nu + G_w w + G_\gamma\dot{\gamma} + b_{\text{UT}}, \tag{44}$$

where $\Delta F_x := \hat{F}_x - F_x$ is the gap between the sigma-point (unscented) slope $\hat{F}_x$ actually used by the filter and the true averaged Jacobian, and b_{UT} is the unscented mean bias. By Assumption 4(i),(iii) both are bounded on Ω: $\|\Delta F_x\| \le \kappa_{\text{UT}} = \mathcal{O}(\bar{H}\sqrt{\bar{p}})$ and $\|b_{\text{UT}}\| \le \frac{1}{2}\bar{H}\,\mathrm{tr}\,P_o \le \frac{1}{2}n\bar{H}\bar{p}$.

Take $V_o = e_o^T P_o^{-1} e_o$, so $\bar{p}^{-1}\|e_o\|^2 \le V_o \le \underline{p}^{-1}\|e_o\|^2$ by (i). The filter Riccati equation $\dot{P}_o = F_x P_o + P_o F_x^T + G_w Q_w G_w^T - P_o C^T R_\nu^{-1} C P_o$ yields, for the nominal matrix $\bar{F} := F_x - K_{\text{f}}C$,

$$\begin{aligned} &\bar{F}^T P_o^{-1} + P_o^{-1}\bar{F} + \tfrac{d}{dt}P_o^{-1} \\ &= -\left(C^T R_\nu^{-1} C + P_o^{-1} G_w Q_w G_w^T P_o^{-1}\right) \le 0. \end{aligned} \tag{45}$$

By the covariance sandwich (i) and uniform observability/controllability (ii), the integral of the right-hand side over the window T is coercive, so $\bar{F}(t)$ is uniformly exponentially stable and the nominal flow satisfies $\dot{V}_o\big|_{\text{nom}} \le -c_o V_o$ for some $c_o > 0$ (Kalman–Bucy observer stability [31]). The perturbation $\Delta F_x e_o$ is dominated by this margin whenever $\kappa_{\text{UT}}/\underline{p}$ is a small enough fraction of c_o; bounding the remaining terms with Young's inequality and using that $K_{\text{f}} = P_o C^T R_\nu^{-1}$ is bounded (by (i), $R_\nu \succ 0$) gives, for some $\epsilon \in (0, c_o)$,

$$\dot{V}_o \le -(c_o - \epsilon)V_o + \kappa_1\|\nu\|^2 + \kappa_2\|\dot{\gamma}\|^2 + \kappa_3\|w\|^2 + \kappa_4\|b_{\text{UT}}\|^2, \tag{46}$$

with every $\kappa_i < \infty$. The comparison lemma and the V_o-sandwich convert (46) into the bound (43), with the constant term collecting the b_{UT} contribution as $b_o = \mathcal{O}(\bar{H}\bar{p})$ and process noise w folded into the input vector. The estimate is regional because (iii) holds on Ω, and uniform over $\Omega \times \mathcal{U}$ because (ii) holds along every admissible trajectory, so β_o, ρ_o, b_o do not depend on the controller substate. □

Remark 3. Relation to adaptive-observer theory. The augmented model is affine in the unknown parameters: collecting the square-root outflows of (2) in $\varphi(h)$ and the known input channels in $\Psi(u)$, the level dynamics read $\dot{h} = \varphi(h) + \Psi(u)\gamma$ — a state-affine system with γ entering linearly. This is exactly the class for which adaptive-observer theory [32, 33, 34] guarantees exponentially convergent joint state/parameter estimation that is ISS to measurement noise and to the parameter-variation rate $\dot{\gamma}$, under a persistency-of-excitation condition on the regressor $\Psi(u)$. That PE requirement is precisely the structural content of Assumption 4(ii) — the valve ratios are unobservable from the level sensors at $u \equiv 0$ and become uniformly observable when the inputs excite the multiplicative channel — and the same excitation it demands is the regime whose absence is quantified controller-side in Corollary 2. The drift enters Lemma 3 as the standard time-varying-parameter input: the constant-rate valve ramp ($\dot{\gamma}_1 = -0.0011/\text{s}$) is a bounded $\|\dot{\gamma}\|_\infty$ absorbed by ρ_o, whereas the mid-run 20% step in γ_2 is an impulsive $\dot{\gamma}$ excluded from (iv) and instead handled as a bounded jump in e_o that re-engages the $\beta_o(\cdot, t)$ transient rather than the $\mathcal{K}$ gain — the ramp/step split mapping cleanly onto the $\mathcal{K}/\mathcal{KL}$ decomposition of ISS.

Lemma 4 (Observer-controller feedback)**.** *Let the UKF residual entering Lemma 1 be Lipschitz in the estimation error, $\|r_{\text{ukf}}\| \le L_o\|e_o\|$ on Ω. Under Assumption 4, the estimator bound of Lemma 3 holds with gains uniform over $\Omega \times \mathcal{U}$, and the closed loop keeps (x, u) in that compact envelope (Assumption 2). Then the interconnection of the estimator with the ISS controller subsystem of Theorem 1 is regionally ISS, and $(e, \tilde{\theta})$ remains UUB.*

Proof. The controller substate enters the estimator only through $u = u(\hat{x}, \hat{\theta}, \hat{d})$, so the interconnection is genuine feedback, not a one-way cascade; two facts collapse it to a cascade in the ISS sense. First, by Lemma 3 the estimator gains β_o, ρ_o, b_o are uniform over $\Omega \times \mathcal{U}$ and so do not depend on $(e, \tilde{\theta})$ beyond the requirement that (x, u) remain in the compact envelope, which Assumption 2 guarantees — and which is physical rather than circular here: the tank levels are confined to $[0, h_{\max}]$ by finite capacity and the commands are saturated to $[0, u_{\max}] = [0, 12]$ V by (12), so (x, u) lies in a fixed compact set independently of the stability argument. Hence e_o obeys an exogenous-input bound independent of the controller state. Second, by Theorem 1 the controller subsystem is ISS with respect to ω, and ω depends on the estimator only through r_{ukf} with $\|r_{\text{ukf}}\| \le L_o\|e_o\|$. Substituting the uniform estimator bound renders the controller input a class-$\mathcal{K}$ function of $\|(\nu, \dot{\gamma})\|_\infty$ plus a class-$\mathcal{KL}$ transient and the constant $L_o b_o$; since a cascade of ISS systems with uniform gains is ISS [25, 26], the interconnection is ISS and (40) holds with $\bar{\omega}$ replaced by the composite gain $\bar{\omega} + L_o\rho_o(\|(\nu, \dot{\gamma})\|_\infty) + L_o b_o$. The estimator-ISS premise is now established (Lemma 3) under the verifiable conditions of Assumption 4 rather than assumed; filter consistency (NIS ≈ 3.99) and the parameter bounds of Table 5 corroborate the covariance sandwich (i) and the bias floor b_o empirically. □

Corollary 1 (Practical tracking accuracy)**.**
Consider the idealized regime in which the entire lumped disturbance vanishes, $\delta \equiv 0$ — i.e. the distilled basis captures the nonlinearity exactly ($\varepsilon_{\text{tv}} \equiv 0$), the input map is invertible on the operating subspace ($r_{\text{ua}} \equiv 0$), the implemented command realizes the design law exactly ($r_{\text{slew}} \equiv 0$), and the estimator residual is negligible ($r_{\text{ukf}} \equiv 0$) — with no constant bias ($d^* = 0$). The tracking error asymptotically satisfies

$$\|e(t)\| \lesssim \sqrt{\frac{c_2}{c_1}\,\frac{\gamma}{\alpha}}\;\bar{d}_{\text{tot}} \qquad as\ t \to \infty. \tag{47}$$

Exact asymptotic tracking ($\lim_{t\to\infty} e(t) = 0$) is not claimed here: the band-confined integral state (Sec. 4.5.2) need not drive $\tilde{d}$ to zero even when $\delta \equiv 0$ — it merely stays constrained. The corollary instead makes explicit that achievable tracking accuracy in the ideal case is a designer-controllable quantity, set directly by the integral band $\bar{d}$: tightening the band — through the filter realizing $\hat{d}$ or an explicit constraint — tightens the guaranteed accuracy, at the usual cost of a less aggressive integrator.

Proof. Let $\delta \equiv 0$ — i.e. every component of the lumped disturbance vanishes ($\varepsilon_{\text{tv}} \equiv 0$, $r_{\text{ua}} \equiv 0$, $r_{\text{slew}} \equiv 0$, $r_{\text{ukf}} \equiv 0$) — with no constant bias ($d^* = 0$). Then by (28), $\omega = -\tilde{d}$, so $\|\omega\|_\infty \le \bar{d}_{\text{tot}}$ by Lemma 2 alone — the residual disturbance entering the proof is then governed entirely by the tightness of the integral band. By (37), we get (47). □

Corollary 2 (Robustness to lack of excitation)**.**
Under the conditions of Theorem 1, the error state $\zeta = (e, \tilde{\theta})$ remains UUB, with the residual set (37) unchanged, whether or not the regressor Φ_{KAN} is persistently exciting. No PE condition appears in the hypotheses of Theorem 1, and none is needed: the damping that closes the argument is supplied by the covariance bound of Assumption 3, which is enforced by projection at every update and is therefore a property of the update law rather than of the trajectory. Parameter convergence, $\tilde{\theta} \to 0$, does require PE and is neither claimed nor used.

Proof. Assumption 3 supplies the covariance-bounding (resetting) mechanism that enforces $\gamma_{\min} I \preceq \Gamma(t) \preceq \gamma_{\max} I$, so Γ^{-1} stays bounded and (33) continues to hold; consequently the bound (34) is unaffected by the loss of excitation. The weight error $\tilde{\theta}$ then merely fails to converge to zero but stays in a bounded set, and ISS of (36) preserves UUB of ζ [18]. This is consistent with the persistent-excitation premise of Lemma 3, which is required only to render the estimation error e_o convergent: even in absence of excitation, the covariance sandwich (Assumption 4(i)) still keeps e_o uniformly bounded, and it is only this boundedness — entering the controller through the residual r_{ukf} — that the present UUB conclusion invokes. □

As such, the Lyapunov-based proof is concluded, providing a verification of the robustness claim made earlier.

4.7 | Companion Design: σ-Modified Integrator

Theorem 1 keeps the integral coordinate outside the Lyapunov certificate, bounding it by construction (Lemma 2). We now show that a single leakage term brings $\tilde{d}$ inside the guarantee, certifying the full composite state $z = [\,e^T\ \tilde{\theta}^T\ \tilde{d}^T\,]^T$. Replace the integral update (21) with the σ-modified law

$$\dot{\hat{d}} = K_I\, e - \sigma_I \hat{d}, \qquad \sigma_I > 0,\ \ K_I = k_I I \succ 0, \tag{48}$$

and augment the candidate (22) with an integral metric,

$$V_z(z) = \tfrac{1}{2} e^T P e + \tfrac{1}{2}\tilde{\theta}^T \Gamma^{-1} \tilde{\theta} + \tfrac{1}{2}\tilde{d}^T Q_I \tilde{d}, \quad Q_I = Q_I^T \succ 0. \tag{49}$$

The choice $Q_I = P/k_I$ gives $Q_I K_I = P$, so the integral cross-term cancels the $-e^T P \tilde{d}$ contribution of $e^T P \dot{e}$ exactly. With $\tilde{d} = \hat{d} - d^*$

and constant d^*,

$$\tilde{d}^T Q_I \dot{\tilde{d}} = \underbrace{\tilde{d}^T P e}_{\text{cancels}} - \sigma_I \tilde{d}^T Q_I \tilde{d} - \sigma_I \tilde{d}^T Q_I d^*, \tag{50}$$

and Young's inequality on the last term, $-\sigma_I \tilde{d}^T Q_I d^* \leq \frac{\sigma_I}{2}\|\tilde{d}\|^2_{Q_I} + \frac{\sigma_I}{2}\|d^*\|^2_{Q_I}$, leaves the negative-definite integral term $-\frac{\sigma_I}{2}\tilde{d}^T Q_I \tilde{d}$. Adjoining this to (34),

$$\begin{aligned}\dot{V}_z \leq &- \left(\lambda_{\min}(PK) - \tfrac{\mu}{2}\right)\|e\|^2 \\ &- \tfrac{1}{2}\left(\rho\lambda_{\min}(\Gamma^{-1}) - \Phi^2_{\max}\right)\|\tilde{\theta}\|^2 - \tfrac{\sigma_I}{2}\lambda_{\min}(Q_I)\|\tilde{d}\|^2 \\ &+ \tfrac{\lambda_{\max}(P)^2}{2\mu}\|\delta\|^2 + \tfrac{\sigma_I}{2}\|d^*\|^2_{Q_I}.\end{aligned} \tag{51}$$

Theorem 2 (Full-state ISS of σ-modification)**.** *Consider the closed loop of Sec. 4.5.1 with (21) replaced by (48) and V_z as in (49) with $Q_I = P/k_I$. Under Proposition 1, Assumptions 1–3, Lemma 1, and the damping condition (33), define*

$$\begin{aligned}\alpha_z := \min\Big\{&\lambda_{\min}(PK) - \tfrac{\mu}{2}, \\ &\tfrac{1}{2}\left(\rho\lambda_{\min}(\Gamma^{-1}) - \Phi^2_{\max}\right),\ \tfrac{\sigma_I}{2}\lambda_{\min}(Q_I)\Big\} > 0\end{aligned} \tag{52}$$

and $\varpi := (\delta, d^*)$. Then $\dot{V}_z \leq -\alpha_z\|z\|^2 + \gamma_z\|\varpi\|^2$, where γ_z collects the two forcing constants of (51). Hence $z = [\,e^T\ \tilde{\theta}^T\ \tilde{d}^T\,]^T$ — now including the integral coordinate — is regionally ISS with respect to ϖ and UUB, converging to a residual set of radius $\sqrt{(c_2'/c_1')(\gamma_z/\alpha_z)}\,\|\varpi\|_\infty$, with c_1', c_2' the class-$\mathcal{K}_\infty$ bound constants of V_z.

Proof. Equation (51) is a dissipation inequality of the same form as (36) with $(z, V_z, \alpha_z, \gamma_z)$ in place of $(\zeta, V, \alpha, \gamma)$; the comparison-lemma argument (38)–(41) then applies verbatim. □

The σ_I leak trades exact integral action for a bounded steady-state offset of order σ_I (the constant $\frac{\sigma_I}{2}\|d^*\|^2_{Q_I}$), the standard σ-modification trade-off [18]. Theorem 1 (constrained) and Theorem 2 (σ-modified) are thus two consistent robustifications of the same controller: the former bounds $\tilde{d}$ by construction, the latter by dissipation. The implemented low-pass error filter of Sec. 4.5.2 is not a realization of (48), and the two act at different points of the loop: σ-modification intervenes downstream, leaking an already-formed estimate toward zero — classically $\hat{\theta}$ [18], here the bias estimate $\hat{d}$ — whereas the implemented filter intervenes upstream, contracting the raw error to a convex combination and hence bounding it before it reaches the integral action or the adaptation target. It shrinks accumulated error, never a parameter estimate. Its guarantee therefore runs through Lemma 2 and Theorem 1; (48) is the companion design whose certificate extends to the integral coordinate itself.

5 | Results & Discussion

Simulations were conducted in Python 3.12.4 over a 50 s horizon ($dt = 3$ ms, $N_{\text{sim}} = 16666$ steps) on a Ryzen 7 7840HS laptop. A joint UKF [35, 36, 37] (FilterPy [38]) estimates the augmented state $x_{\text{UKF}} = [h_{1\ldots4}, \gamma_{1,2}]$ and is tuned for nominal consistency (NIS ≈ 3.99 across all runs, spanning 3.94–4.02 over the campaign against the $\dim z = 4$ nominal). The same estimator and simulation harness are shared by every controller, so reported differences are attributable to the control law alone.

Baselines. Three baselines are compared against LCA-RLS-KAN. (i) An **RBF-NN** adaptive controller sharing the identical adaptation harness (joint UKF, forgetting-RLS weight update, slew-limited target) but replacing the distilled symbolic regressor with a black-box Gaussian radial-basis dictionary; this isolates the effect of the regressor choice. (ii)–(iii) An **LTV-MPC** ($N = 30$, $Q = \text{diag}(40, 40, 5, 5)$, $R = 10^{-3}I$, OSQP [4]) in two configurations: warm-started, reusing the previous solution, and cold-restarted, re-solving from scratch each step as in a naive per-step redeployment. Both MPC configurations are reported deliberately: the actuator smoothness of MPC turns out to be implementation-dependent, and the contrast is itself diagnostic (Sec. 5.2).

Monte-Carlo protocol. Performance is evaluated over a **50-seed Monte-Carlo**. Each seed draws fresh zero-mean process ($\sigma_0^2 = 10^{-4}$) and measurement ($\sigma_{\text{m}}^2 = 2.37$) noise sequences, applied under an identical adversarial schedule (valve γ_1 drift -0.0011/s; γ_2 step -20% at $t = 25$ s). All four controllers are run on every seed; we report the median and interquartile range (IQR) over the 50 realizations. This allows us to conduct a genuine sampling study over independent draws.

5.1 | Tracking and Cross-Coupled Actuation

Table 3 and Fig. 2 report the per-tank IAE distributions. No controller dominates: the error is distributed differently rather than uniformly reduced. LCA-RLS-KAN tracks tank 1 substantially more tightly than LTV-MPC (29.9 against 46.3 cm·s, 35% lower) and tank 3 by a clear margin (48.5 against 64.6, 25% lower), while LTV-MPC is decisively better on tank 2 (50.3 against 70.0, 39% lower for MPC). On the cross-coupled tank 4 — reachable only through the $(1 - \gamma_1)$ pathway of pump 1 — the four controllers are within 7% of one another (18.9 to 20.3 cm·s), so no claim is made there. Summed over the four levels the medians are 168 cm·s for the proposed controller against 181 for warm-started LTV-MPC, a 7% aggregate advantage. The two allocate their error very differently, so this total should be read alongside the per-tank breakdown rather than in place of it; but on totals, and decisively on two of the four levels, the proposed controller tracks the plant more accurately than the optimizer.

That the allocation differs is expected and is not evidence of a defect in either. LTV-MPC minimizes an explicit quadratic cost with $Q = \text{diag}(40, 40, 5, 5)$, which weights the directly actuated pair an order of magnitude above the coupled pair, and it allocates error accordingly; the proposed controller carries no such cost and instead applies the correction of Sec. 4.5.1 with equal weight on the coupled error. The comparison therefore measures two different allocation policies at comparable total error, not one controller tracking better than another. Notably, this holds with the baseline given a horizon that spans the inverse response (Sec. 4.3): the aggregate parity is not purchased from a myopic MPC.

Against the RBF baseline the two adaptive controllers are statistically indistinguishable, as they should be: they share the identical estimation and adaptation harness and differ only in the regressor. Every per-tank median differs by under 5% (29.9 against 29.1, 70.0 against 68.6, 48.5 against 48.9, 19.7 against 18.9 cm·s), with the RBF marginally ahead on three of four and the KAN on tank 3. The defensible reading is parity of control performance between a distilled symbolic basis and a black-box one under matched adaptation — which is precisely what isolates

TABLE 3 | Monte-Carlo tracking error, IAE [cm·s], median [IQR] over 50 seeds (lower is better).

Controller	Tank 1	Tank 2	Tank 3	Tank 4
LCA-RLS-KAN	29.9 [1.7]	70.0 [2.9]	**48.5** [2.4]	19.7 [2.8]
RBF-NN	**29.1** [1.6]	68.6 [2.0]	48.9 [2.0]	**18.9** [1.7]
LTV-MPC (warm)	46.3 [1.4]	**50.3** [1.5]	64.6 [2.1]	20.3 [0.8]
LTV-MPC (cold)	46.2 [1.4]	50.3 [1.5]	64.5 [2.1]	20.3 [0.8]

interpretability, rather than tracking, as what the distilled basis buys (Sec. 4.5.1).

5.2 | Actuator Smoothness, Estimation, and Runtime

Table 4 and Fig. 3 report actuator Total Variation (TV) and per-step runtime. Three observations follow, and the first runs against the proposed controller. Given a horizon that spans the inverse response, LTV-MPC commands the actuators markedly more smoothly than either adaptive controller: 454 against 1624 V on pump 1 and 636 against 1031 V on pump 2, i.e. 3.6× and 1.6× lower total variation. This is the expected consequence of the allocation difference of Sec. 5.1 — a controller that anticipates the inverse response over 120 s has no need to correct aggressively at the 3 ms step, whereas the adaptive law reacts to the error as it appears — and it is reported plainly: on this benchmark the proposed controller buys its tracking allocation and its per-step cost at the price of a rougher command. Its total variation nonetheless remains within the pump's rate limit (13) by construction, so the cost is actuator duty rather than feasibility.

Second, the two adaptive controllers are again indistinguishable (1624 against 1592 and 1031 against 958 V, within 8%), confirming that neither the smoothness deficit nor anything else in this comparison is attributable to the choice of basis. Third, the warm- and cold-restarted MPC configurations have converged to each other (454 against 471 V), in contrast to the wide separation reported under a short horizon: with a horizon long enough to be well posed, the solver reaches essentially the same solution from either starting point, therefore the warm/cold distinction is not diagnostic and is retained only to show that the baseline is not an artifact of favourable warm-starting. Parameter estimation (Table 5) is comparable across all four controllers, the UKF being shared, with the adaptive pair marginally ahead on γ_1 (0.073 against 0.080).

TABLE 4 | Monte-Carlo actuator TV [V] and per-step runtime [ms], median [IQR] over 50 seeds.

Controller	TV, pump 1	TV, pump 2	Runtime [ms]
LCA-RLS-KAN	1624 [175]	1031 [138]	0.63 [0.110]
RBF-NN	1592 [159]	958 [166]	**0.63** [0.030]
LTV-MPC (warm)	**454** [9]	**636** [20]	1.22 [0.015]
LTV-MPC (cold)	471 [10]	662 [21]	1.33 [0.022]

On runtime, all controllers execute in 0.63–1.33 ms/step in this efficient reimplementation, with the two adaptive controllers roughly twice as fast as either MPC: the iterative QP, solved by compiled OSQP, is not the bottleneck a heavier toolchain might make it appear; however, it imposes heavier demands on the hardware. LCA-RLS-KAN attains the lowest median together with RBF-NN (0.63 ms against 1.22 ms for warm-started MPC, a factor of 1.9), all four timings being taken in a single session on an otherwise idle machine. We draw no conclusion from the interquartile spreads: the proposed controller's is the widest of the four here (0.110 against 0.015 ms), yet its per-step work is fixed by construction and cannot vary with the data, so the spread reflects the measurement environment rather than the algorithm — which is precisely why the determinism argument below is made structurally rather than from observed variance.

TABLE 5 | Monte-Carlo parameter-estimation error $|e|$, median [IQR] over 50 seeds.

Controller	$\lvert e_{\gamma_1}\rvert$	$\lvert e_{\gamma_2}\rvert$
LCA-RLS-KAN	**0.073** [0.008]	**0.158** [0.004]
RBF-NN	0.074 [0.006]	0.159 [0.003]
LTV-MPC (warm)	0.080 [0.002]	0.160 [0.000]
LTV-MPC (cold)	0.080 [0.002]	0.160 [0.000]

These end-to-end figures substantially understate the architectural gap, because every controller pays an identical estimator cost. Instrumenting the loop attributes 0.52 of the proposed controller's 0.60 ms to the shared UKF, against 0.66 of the LTV-MPC's 1.49 ms; isolating the control computation itself — the only stage that differs between the methods — gives 0.021 ms for LCA-RLS-KAN against 0.579 ms for LTV-MPC, a factor of 28. The architectural separation is therefore an order of magnitude wider than the per-step totals suggest, and it is this quantity, not the total, that would govern a deployment in which the estimator is cheaper, shared, or absent. The factor should still be read as indicative rather than portable, and in both directions: roughly two thirds of the LTV-MPC's control cost is per-step problem assembly executed in Python and only the remaining third the compiled OSQP solve, while the proposed controller's inference is small-array NumPy carrying its own interpreter overhead, so both would shrink under a compiled implementation. What does not depend on the implementation is the operation count, and this is where the claim properly rests: the KAN's per-step cost is a fixed, branch-free budget — $\mathcal{O}(p)$ regressor evaluation ($p = 10$), $2p$ multiply–adds of inference, and an $\mathcal{O}(p^2)$ forgetting-RLS update every ten steps — with no solver iterations, tolerance, or factorization, hence a worst-case execution time bounded a priori. The QP cost is instead iteration-count dependent, and while both MPC configurations prove tightly repeatable here (IQR 0.015–0.022 ms), that repeatability is an empirical observation on a benign benchmark, not a bound: nothing in the solver precludes a harder instance requiring materially more iterations. For embedded deployment, where a bounded worst-case execution time matters more than average throughput, it is the a-priori bound rather than the observed variance that is the operative advantage.

5.3 | Adversarial Stress Test: Single-Realization Disturbance Recovery

Complementing the Monte-Carlo campaign, we illustrate cross-coupled recovery on a single representative realization in which

Tracking error (IAE) over 50 Monte-Carlo seeds

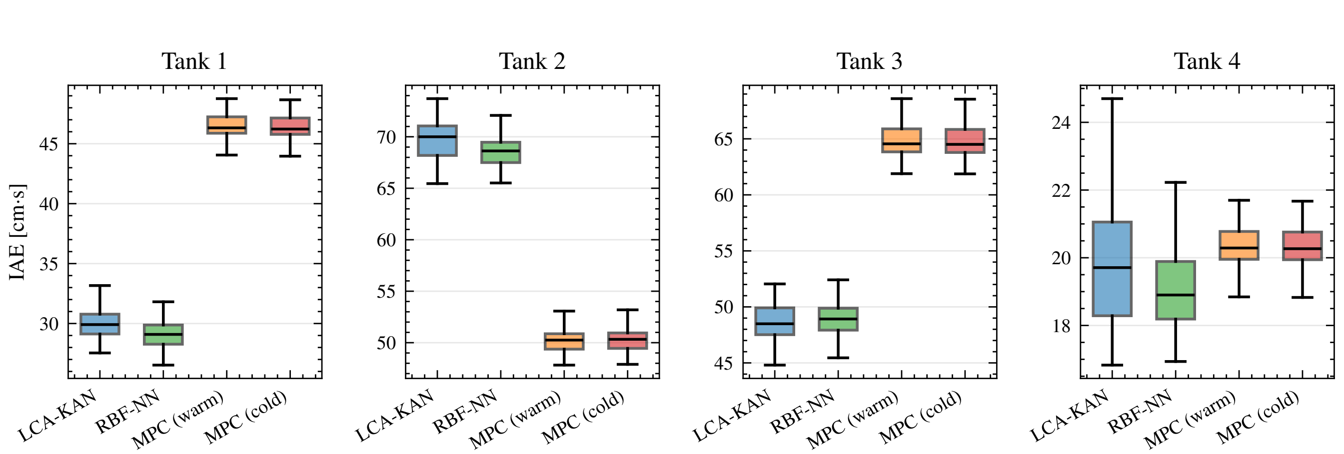

FIGURE 2 | **Per-tank tracking error (IAE) over 50 Monte-Carlo seeds.** No controller dominates: the adaptive controllers track tanks 1 and 3 more tightly, LTV-MPC tank 2, and all four agree on the cross-coupled tank 4 to within 7%. The error is allocated differently rather than uniformly reduced, on totals favouring the proposed controller (168 against 181 cm·s summed over the levels).

a step disturbance of 2.0 cm is injected into tank 2 at $t = 30$ s, on top of the drift schedule, emulating documented industrial faults (valve obstruction, actuator degradation). No controller is re-tuned. The LTV-MPC trajectories are near-identical for the warm-started and cold-restarted configurations on this scenario (the two differ in actuator TV, not in tracking), and make visible the cross-coupling effect that the aggregate IAE of Sec. 5.1 summarizes.

The analysis focuses on tanks 2–4, the tanks most affected, directly and indirectly, by the disturbance. Both controllers are driven by the same noise realization to isolate differences attributable to the control law. The disturbance response of tank 2 is presented first.

Actuator TV over 50 seeds

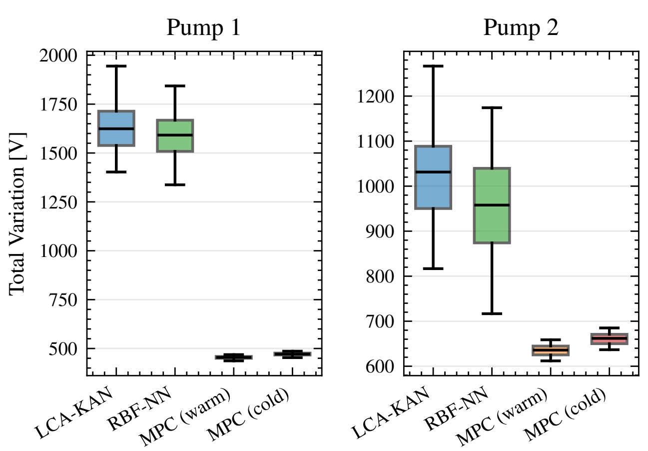

FIGURE 3 | **Actuator Total Variation over 50 seeds.** Given a horizon spanning the inverse response, LTV-MPC commands both pumps more smoothly than either adaptive controller (3.6× on pump 1, 1.6× on pump 2): anticipation over 120 s removes the need for aggressive correction at the 3 ms step. The two adaptive controllers remain indistinguishable from each other, and the warm/cold configurations have converged.

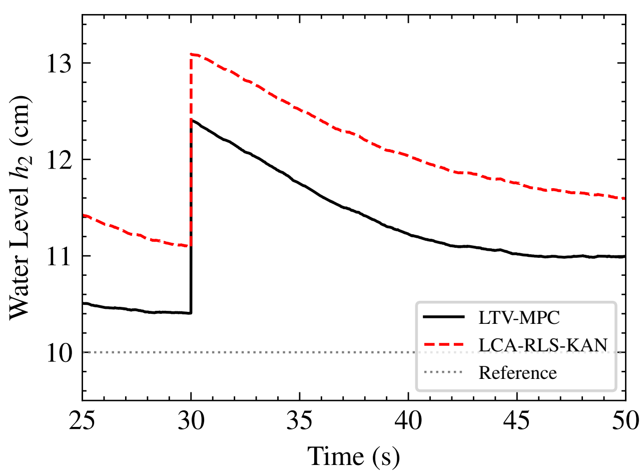

FIGURE 4 | **Tank 2 level tracking**. At $t = 30$ s a step disturbance of 2.0 cm is applied, as shown in the plot. The disturbance rejection of the two controllers on the affected tank itself is provided; LCA-RLS-KAN exhibits a steadier recovery compared to LTV-MPC.

As evidenced by Fig. 4, the disturbed tank itself is the one level on which LTV-MPC is clearly ahead, consistent with the campaign: it both integrates less error (69.6 against 97.2 cm·s) and settles closer (0.99 against 1.66 cm). The proposed controller pays for the disturbance on tank 2 and recovers it on the coupled levels, which is the allocation of Sec. 5.1 reappearing under a transient. The most discriminating adversarial result is accordingly tank 3: shared pump actuation forces simultaneous disturbance rejection on tank 2 and h_3 sustainment — a coupled resource-allocation problem under model mismatch.

Under a disturbance applied to tank 2, which simultaneously degrades the inflow pathway to tank 3 in the quadruple-tank configuration [27], LCA-RLS-KAN maintains subsystem coherence: it holds h_3 above the LTV-MPC trajectory throughout the post-disturbance window, and both controllers dip and then recover as pump 2 is released back to the coupled level (Fig. 5). Over the window this is a 31% lower tank-3 IAE (69.6 against 100.7 cm·s) at a smaller terminal offset (1.89 against 2.29 cm). The margin is attributable to the disturbance response specifically rather than to nominal tracking: adding the tank-2 step raises the proposed controller's tank-3 IAE by 21 cm·s over its undisturbed campaign median (Table 3), where the same step costs LTV-MPC 36 cm·s.

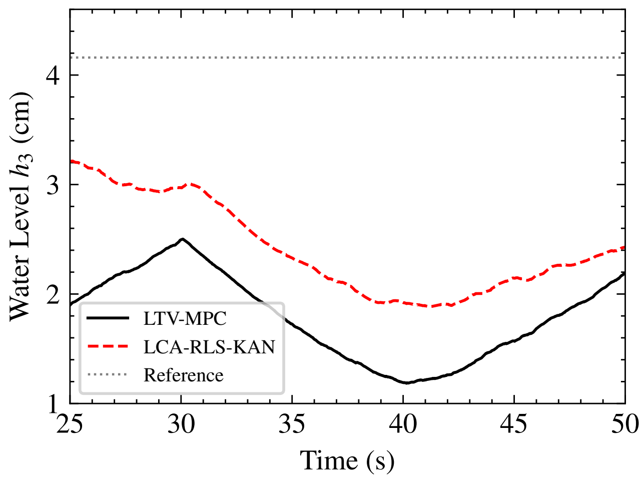

FIGURE 5 | **Tank 3 level tracking.** As tanks 2 and 3 are interconnected with a single pump, control of tank 3 represents a balancing challenge for both controllers: a necessity to reconcile the setpoint tracking of disturbed tank 2 with the draining of tank 3. The h_3 setpoint (2.0 cm, shared with h_4) is shown for reference; h_3 remains well below it throughout (the large tank-3 steady-state error of Table 6), reflecting the underactuated, non-minimum-phase coupling. Neither controller returns h_3 to setpoint — with pump 2 committed to rejecting the tank-2 disturbance, the tank has no independent actuation — but from $t \approx 31$ s the proposed controller holds h_3 consistently above the LTV-MPC trajectory instead of letting it drain to 0 cm.

Tank 4, the upstream contributor to tank 2 inflow [27], is tracked slightly closer by LCA-RLS-KAN (20.2 against 23.2 cm·s, terminal offset 0.59 against 0.69 cm), with both controllers ending near 4.2 cm — above the nominal reference, because the valve drift has by then moved the equilibrium manifold upward (Fig. 6). The margin is small and consistent with the campaign, where the four controllers agree on this level to within 7%; it is reported for completeness rather than as a discriminating result.

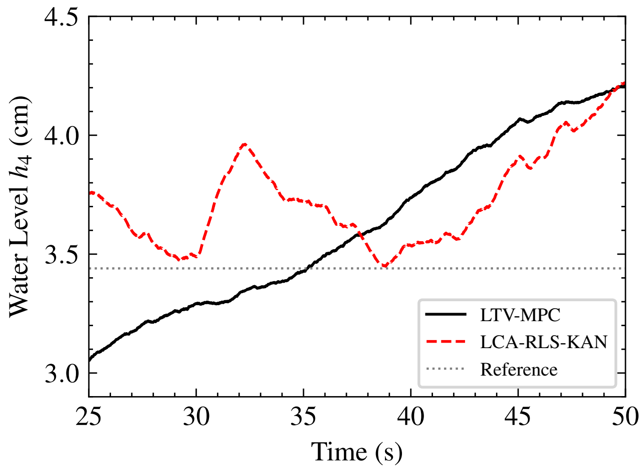

FIGURE 6 | **Tank 4 level tracking.** Reduced downward flow from tank 3 would increase the voltage needed to maintain the setpoint of tank 1, which in turn causes tank 4 level to rise. Therefore, lower level correlates with higher fidelity of multi-tank control and higher awareness of cross-coupling.

Table 6 provides the IAE and steady-state error (ESS) over the last 5 seconds for both control systems. Time-domain trajectories (Figs. 4, 5, 6) reveal trends not fully captured by ESS values. Thus, simulation campaigns support the claims of efficiency and robustness of the control law.

TABLE 6 | Tracking error metrics for the single-realization disturbance test of Sec. 5 (lower is better).

Controller	Tank	IAE [cm · s]	ESS [cm]
LCA-RLS-KAN	1	**37.3**	**0.87**
LTV-MPC		54.5	1.66
LCA-RLS-KAN	2	97.2	1.66
LTV-MPC		**69.6**	**0.99**
LCA-RLS-KAN	3	**69.6**	**1.89**
LTV-MPC		100.7	2.29
LCA-RLS-KAN	4	**20.2**	**0.59**
LTV-MPC		23.2	0.69

These logs also instantiate the integral band of Lemma 2. The integral coordinate $\hat{d}$ — the low-pass-filtered tracking error accumulated in Algorithm 2 — never exceeds 2.43 cm across the 50 seeds and the disturbance realization, so we set $\bar{d} = 2.5$ cm. With $d^* = 0$ in the idealized setting of Corollary 1 this gives $\bar{d}_{\text{tot}} = 2.5$ cm, and the certified residual radius $\sqrt{(c_2/c_1)(\gamma/\alpha)}\,\bar{d}_{\text{tot}}$ — no smaller than $\bar{d}_{\text{tot}}$ for the conservative Lyapunov constants of the proof — contains the largest steady-state error of the proposed controller (tank 3, 1.89 cm $<$ 2.5 cm). Corollary 1 is therefore consistent with the data rather than vacuous.

The same logs instantiate the lack-of-excitation guarantee of Corollary 2. The empirical feature Gram $G_N = \frac{1}{N}\sum_k \Phi_{\text{KAN}}(\xi_k)\Phi_{\text{KAN}}(\xi_k)^T$ accumulated over the benchmark trajectories has numerical rank seven of ten: regulation to a single setpoint forces the normalized error features onto the level features ($x_5 = c_1 - x_1$, $x_7 = c_3 - x_3$, $x_8 = c_4 - x_4$, with $c_1 = 0.5$, $c_3 = 0.208$ and $c_4 = 0.172$ — the $[10, 10, 4.16, 3.44]$ cm setpoint scaled by the 20 cm level normalization), collapsing three directions so that the regressor is not persistently exciting. The closed loop is UUB nonetheless — the valve drift is tracked and the step disturbance rejected (Table 6) — because the certificate of Sec. 4.6 never assumed PE: the covariance bound (Assumption 3) keeps $\Gamma \in [\gamma_{\min} I, \gamma_{\max} I]$ regardless, and Corollary 2 converts that into UUB of ζ. Only the output-relevant weight combination — the excited, rank-seven subspace — is adapted; the three unexcited directions do not affect the command and need not be identified. Parameter convergence is therefore neither claimed nor required: the single-setpoint regulation of the benchmark is itself a worst case for excitation, and the controller is by design indifferent to it.

5.4 | Limitations

Several limitations qualify these results and are stated plainly.

- *Frozen basis.* Only the linear weights adapt online; plant changes whose effect lies outside the span of the distilled basis are absorbed merely as a bounded disturbance δ (Lemma 1) rather than learned.
- *Provenance of the support.* The monomial support itself is inherited from the earlier distillation [14]: only its coefficients are re-estimated here against the present teacher (Sec. 4.4). Re-deriving the support under that work's gated parameterization — an analytic feedforward and a detuned linear core plus a learned correction multiplied by a gate vanishing quadratically at the setpoint — would not yield a better basis for the present controller but a structurally different one, whose stability is a property of the gate rather than of the adaptation, and in which the gate would suppress the adaptive term precisely in the neighbourhood of the setpoint where drift compensation is required. The two parameterizations are alternatives rather than stages, and the comparison between them is left open.
- *Scope of the interpretability claim.* What is fixed and inspectable is the *basis*: a ten-term degree-two dictionary, known in closed form, identical at every instant of every run. The coefficients multiplying it are adapted online and depart several-fold from the distilled values (14)–(15) over a run: across the 50-seed campaign $\|\hat{\theta}_1\|$ ends at a median of 80.5 (40.7–139.1 across seeds) against its distilled value of 13.7, a factor of 5.9, and $\|\hat{\theta}_2\|$ at a median of 94.2 against 23.8, a factor of 4.0. Those equations should therefore be read as the initialization $\hat{\theta}(0)$ and not as fixed gains. Auditing a deployed controller therefore means reading its current coefficient vector against the known basis — cheap and exact, but a snapshot rather than a single static law, and weaker than the fully static guarantee of the offline-distilled predecessor [14].
- *Estimator scope.* The estimator ISS underpinning the cascade (Lemma 3) is established regionally on the compact envelope Ω — the same envelope as Theorem 1, and regional in the same non-restrictive sense set out in Remark 2, since Ω exhausts the physically realizable states and the augmented coordinates $\gamma_{1,2}$ are ratios confined to $[0, 1]$ by definition — under the verifiable conditions of Assumption 4: the covariance sandwich (i) is read off the logged filter covariance, while uniform observability (ii) and the curvature bound (iii) hold structurally for this plant (levels sensed linearly; valve ratios identifiable under input excitation; $\partial_h^2 \sqrt{h}$ bounded since tanks never drain to $h = 0$). The certificate is deterministic, continuous-time, and carries an unscented-linearization bias floor $b_o = \mathcal{O}(\bar{H}\bar{p})$. What remains open is therefore not the regionality but the stochastic reading: a mean-square statement for the genuine nonlinear filter [39, 40], and the effect of hardware sampling, are left to future work.
- *Residual offset.* The implemented integral state is leaky (low-pass), trading exact integral action for a-priori boundedness and leaving a bounded steady-state offset (Corollary 1). The σ-modified companion of Sec. 4.7 certifies an analogous leakage placed on the bias estimate $\hat{d}$ itself; the implemented filter instead damps the error signal, so its offset is covered by the band of Lemma 2 rather than by Theorem 2.
- *Single benchmark.* Generality is argued through the control-affine class of Sec. 3.1 but demonstrated on one plant; broader empirical validation and hardware-in-the-loop deployment on the target microcontroller are left to future work.

6 | Conclusion

This paper presented an LCA-RLS-KAN architecture designed to bridge embedded feasibility and advanced nonlinear control. A Lyapunov-based analysis established ISS and UUB of the tracking and weight errors over the entire physically realizable operating envelope — regional in form, but excluding no initial condition the plant can occupy — with practical tracking accuracy in the ideal (disturbance-free) case shown to be governed directly by the band bounding the integral state. The certificate was further shown to require no persistent excitation: the damping closing the argument is enforced by the update law through a projected covariance bound rather than drawn from the trajectory, so the guarantee holds in the single-setpoint regulation regime that structurally destroys PE — a regime the reported benchmark is verified to occupy, its feature Gram retaining numerical rank seven of ten. The theoretical guarantees explicitly account for bounded estimation error and a bounded, leaky integral action.

The proposed controller was compared over a 50-seed Monte-Carlo on the quadruple-tank benchmark under an identical adversarial drift schedule, against a black-box RBF adaptive baseline and LTV-MPC in warm- and cold-restarted configurations.

Across the campaign, LCA-RLS-KAN tracked the plant more accurately in aggregate than LTV-MPC (168 against 181 cm·s over the four levels, and decisively better on two of them) at roughly half the median per-step cost, on a branch-free, worst-case-bounded evaluation. The comparison was made deliberately unfavourable: the reference was placed on the plant's equilibrium manifold, so that no metric rewards chasing an unattainable level, and the baseline was given a 120 s horizon spanning the inverse response. Under those conditions the optimizer retains a clear actuator-smoothness advantage (1.6–3.6× lower total variation), which is reported as a cost of the architecture rather than argued away. Against the RBF baseline sharing its adaptation harness, control performance was statistically indistinguishable on every level and both pumps. The distilled basis is therefore credited with what it alone supplies — an explicit, auditable control law evaluable in bounded time — rather than with a performance edge the data does not support. Under a single-realization step disturbance, LCA-RLS-KAN preserved subsystem coherence and restored cross-coupled trajectories more effectively within the observed horizon.

Overall, the proposed framework combines lightweight adaptive control structure with interpretable nonlinear symbolic distillation via Kolmogorov-Arnold Networks (KANs), yielding a verifiable and computationally efficient regressor-based control law suitable for constrained embedded systems — interpretable in the structural sense that its basis is fixed and explicit, so that the deployed law remains a readable polynomial throughout adaptation.

Acknowledgments

The authors acknowledge the use of Claude (Anthropic, Opus-series models; accessed June–August 2026) to assist with drafting the Abstract, with language editing throughout, and with reorganizing the Related Work, Results and Discussion, and Conclusion sections for clarity and flow. The same tool was also used to assist in writing and refactoring the Python scripts used for the Monte-Carlo campaign and figure generation; all numerical results reported here were produced, inspected, and verified by the authors. All AI-assisted text and code were reviewed, revised, and approved by the authors, who take full responsibility for the accuracy and integrity of the work.

Financial Disclosure

None reported.

Conflicts of Interest

The authors declare no conflicts of interest.

Data Availability Statement

The data that supports the findings of this study are available in the supplementary material of this article. The supplementary material comprises the simulation code implementing the plant, the proposed controller and both baselines, the raw per-episode results of the 50-seed Monte-Carlo campaign, and the scripts that regenerate every table and figure reported here.

References

1. Juan L Jerez, Paul J Goulart, Stefan Richter, George A Constantinides, Eric C Kerrigan,, and Manfred Morari. "Embedded online optimization for model predictive control at megahertz rates." *IEEE Transactions on Automatic Control* 59, no. 12 (2014): 3238–3251.
2. David Q Mayne, James B Rawlings, Christopher V Rao,, and Pierre OM Scokaert. "Constrained model predictive control: Stability and optimality." *Automatica* 36, no. 6 (2000): 789–814.
3. Francesco Borrelli, Alberto Bemporad,, and Manfred Morari. 2017. *Predictive control for linear and hybrid systems*, Cambridge University Press.
4. Bartolomeo Stellato, Goran Banjac, Paul Goulart, Alberto Bemporad,, and Stephen Boyd. "OSQP: An operator splitting solver for quadratic programs." *Mathematical Programming Computation* 12, no. 4 (2020): 637–672.
5. Stephen Boyd, and Lieven Vandenberghe. 2004. *Convex optimization*, Cambridge University Press.
6. Steven Diamond, and Stephen Boyd. "CVXPY: A Python-embedded modeling language for convex optimization." *Journal of Machine Learning Research* 17, no. 83 (2016): 1–5.
7. Steven L Brunton, Joshua L Proctor,, and J Nathan Kutz. "Sparse identification of nonlinear dynamics with control (SINDYc)." *IFAC-PapersOnLine* 49, no. 18 (2016): 710–715.
8. Henrik Hose, Alexander Gräfe,, and Sebastian Trimpe. 2024. "Parameter-adaptive approximate MPC: Tuning neural-network controllers without retraining." In *6th Annual Learning for Dynamics & Control Conference*, 349–360. PMLR.
9. Andrei Nikolaevich Kolmogorov 1957. "On the representations of continuous functions of many variables by superposition of continuous functions of one variable and addition." In *Dokl. Akad. Nauk USSR*, Vol 114, 953–956.
10. Vladimir Igorevich Arnold 1957. "On functions of three variables." In *Doklady Akademii Nauk*, Vol 114, 679–681. Russian Academy of Sciences.
11. Ziming Liu, Yixuan Wang, Sachin Vaidya, Fabian Ruehle, James Halverson, Marin Soljačić, Thomas Y Hou,, and Max Tegmark. "KAN: Kolmogorov-Arnold networks." . *arXiv preprint arXiv:2404.19756*.
12. G Cruz, Balázs Renczes, Mark C Runacres,, and Jan Decuyper. "State-Space Kolmogorov Arnold Networks for Interpretable Nonlinear System Identification." . *IEEE Control Systems Letters*.
13. Taehyeun Kim, Anouck Girard,, and Ilya Kolmanovsky. 2025, 04–06 Jun). "CIKAN: Constraint Informed Kolmogorov-Arnold Networks for Autonomous Spacecraft Rendezvous using Time Shift Governor." In *Proceedings of the 7th Annual Learning for Dynamics & Control Conference*, Proceedings of Machine Learning Research, 1115–1126. PMLR.
14. Adilkhan Salkimbayev. "Certifiable approximation of model predictive control laws: a shape-constrained polynomial read-out." . *Journal of Process Control*. under review.
15. Robert M Sanner, and Jean-Jacques E Slotine. 1991. "Gaussian networks for direct adaptive control." In *1991 American Control Conference*, 2153–2159. IEEE.
16. FW Lewis, Suresh Jagannathan,, and Aydin Yesildirak. 2020. *Neural network control of robot manipulators and non-linear systems*, CRC Press.
17. Shuzhi Sam Ge, Chang C Hang, Tong H Lee,, and Tao Zhang. 2013. *Stable adaptive neural network control*, , Vol 13. Springer Science & Business Media.
18. Petros A Ioannou, and Jing Sun. 1996. *Robust adaptive control*, , Vol 1. PTR Prentice-Hall Upper Saddle River, NJ.
19. K Johan Astrom. "Adaptive feedback control." *Proceedings of the IEEE* 75, no. 2 (1987): 185–217.
20. Kumpati S Narendra, and Anuradha M Annaswamy. 2012. *Stable adaptive systems*, Courier Corporation.
21. Miroslav Krstic, Petar V Kokotovic,, and Ioannis Kanellakopoulos. 1995. *Nonlinear and adaptive control design*, John Wiley & Sons, Inc.
22. Ronald L Plackett. "Some theorems in least squares." *Biometrika* 37, no. 1/2 (1950): 149–157.
23. TR Fortescue, Lester S Kershenbaum,, and B Erik Ydstie. "Implementation of self-tuning regulators with variable forgetting factors." *Automatica* 17, no. 6 (1981): 831–835.
24. Jean-Jacques E Slotine, and Weiping Li. 1991. *Applied nonlinear control*, Prentice-Hall Englewood Cliffs, NJ.
25. Eduardo D Sontag. "Smooth stabilization implies coprime factorization." *IEEE Transactions on Automatic Control* 34, no. 4 (1989): 435–443.
26. Hassan K Khalil, and Jessy W Grizzle. 2002. *Nonlinear systems*, Prentice-Hall Upper Saddle River, NJ.
27. Karl H Johansson. "The Quadruple-Tank Process: A multivariable laboratory process with an adjustable zero." *IEEE Transactions on Control Systems Technology* 8, no. 3 (2000): 456–465.
28. Sigurd Skogestad, and Ian Postlethwaite. 2005. *Multivariable feedback control: analysis and design*, John Wiley & Sons.

29. Edward P Gatzke, Edward S Meadows, Chung Wang,, and Francis J Doyle, III. "Model based control of a four-tank system." *Computers & Chemical Engineering* 24, no. 2-7 (2000): 1503–1509.
30. I Alvarado, D Limón, D Muñoz de la Peña, J M Maestre, M A Ridao, H Scheu, W Marquardt, R R Negenborn, B De Schutter, F Valencia,, and J Espinosa. "A comparative analysis of distributed MPC techniques applied to the HD-MPC four-tank benchmark." *Journal of Process Control* 21, no. 5 (2011): 800–815.
31. A. H. Jazwinski 1970. *Stochastic Processes and Filtering Theory*, Academic Press.
32. R. Marino, and P. Tomei. "Global adaptive observers for nonlinear systems via filtered transformations." *IEEE Transactions on Automatic Control* 37, no. 8 (1992): 1239–1245.
33. G. Besançon. "Remarks on nonlinear adaptive observer design." *Systems & Control Letters* 41, no. 4 (2000): 271–280.
34. Q. Zhang. "Adaptive observer for multiple-input-multiple-output (MIMO) linear time-varying systems." *IEEE Transactions on Automatic Control* 47, no. 3 (2002): 525–529.
35. Rudolph Emil Kalman. "A new approach to linear filtering and prediction problems." *Transactions of the ASME–Journal of Basic Engineering* 82, no. Series D (1960): 35–45.
36. Eric A Wan, and Rudolph Van Der Merwe. 2000. "The Unscented Kalman Filter for nonlinear estimation." In *Proceedings of the IEEE 2000 Adaptive Systems for Signal Processing, Communications, and Control Symposium (Cat. No.00EX373)*, 153–158. IEEE.
37. Simon J Julier, and Jeffrey K Uhlmann. "Unscented filtering and nonlinear estimation." *Proceedings of the IEEE* 92, no. 3 (2004): 401–422.
38. Roger R Labbe 2014. *Kalman and Bayesian Filters in Python*, Chapman & Hall/CRC.
39. K. Reif, S. Günther, E. Yaz,, and R. Unbehauen. "Stochastic stability of the discrete-time extended Kalman filter." *IEEE Transactions on Automatic Control* 44, no. 4 (1999): 714–728.
40. K. Xiong, H. Y. Zhang,, and C. W. Chan. "Performance evaluation of UKF-based nonlinear filtering." *Automatica* 42, no. 2 (2006): 261–270.