Robust Adaptive Kolmogorov-Arnold Neural Control

Adilkhan Salkimbayev and Khawaja S. Haider

***Abstract*— Reliable real-time control of non-minimum phase systems remains a bottleneck for embedded platforms due to the computational burden of online optimization, cumbersome architecture and insufficient stability guarantees in the presence of model mismatch. We propose a Lipschitz-Continuous Adaptive (LCA) architecture using Recursive Least Squares (RLS) to update Kolmogorov-Arnold Networks (KANs) for rapid, robust, and interpretable control. The robustness claim is validated via Lyapunov-based analysis, proving Input-to-State Stability and Global Uniform Ultimate Boundedness. High-fidelity quadruple-tank benchmark under adversarial parameter drift demonstrates superior recovery and efficiency on resource-constrained hardware, where the proposed controller achieves a 5x reduction in actuator Total Variation and 5x reduced runtime on a conservative lower bound relative to Linear Time-Varying Model Predictive Control on comparable multi-tank aggregate tracking performance. The architecture demonstrated in this paper formalizes considerable comprehensibility of the control law. The architecture is therefore particularly suited for resource-constrained real-time embedded platforms (e.g., Cortex-M-class microcontrollers), where real-time control evaluation is critical.**

***Index Terms*— Adaptive control, model predictive control, Kolmogorov-Arnold networks, neuro-symbolic control, real-time control, interpretable neural networks, non-minimum phase systems.**

I. 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.

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 'black-box' nature of standard Multi-Layer Perceptrons (MLPs) [8]. By contrast, Kolmogorov-Arnold Networks (KANs) provide an alternative: by utilizing 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, the feasibility of deploying static KANs for feedback control was established, demonstrating preservation of high fidelity of the control law with interpretable structure, which enabled a forensic audit of KAN's output and relative simplicity of the control law. However, the process relied on offline training, leaving the controller vulnerable to unmodeled dynamics and parameter drift once deployed and providing no formal guarantees of robustness.

In this letter, the Lipschitz-Continuous Adaptive RLS-KAN (LCA-RLS-KAN) is proposed. The contributions of this paper are threefold: first, Recursive Least Squares (RLS) [14] is introduced as an update law for the KAN-distilled polynomials that operates online; second, Lyapunov-based theorem establishing Input-to-State Stability (ISS) and Global Uniform Ultimate Boundedness (GUUB) in the presence of approximation errors is provided with corollaries outlining ideal-case asymptotic convergence and robustness under lack of persistent excitation; third, controller's performance on a high-fidelity simulation is demonstrated, establishing the improved execution time and significantly reduced Total Variation (TV) of input, highlighting architectural advantages of the controller in embedded settings.

II. Problem Formulation

A. Mathematical Model

To compare both 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 non-minimum-phase (NMP) behavior [15]. 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 [15], [16].

The quadruple-tank model follows these equations [15]:

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

A. Salkimbayev and K. S. Haider are with the School of Information Technology and Engineering, Kazakh-British Technical University, Almaty, Kazakhstan (email: a_salkimbayev@kbtu.kz; s.khawaja@kbtu.kz).

In (1), 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. Plant diagram is provided in [17].

B. Control Objective

The plant configuration follows Johansson's NMP setup, described in [15]. The target state is set as $[10.0, 10.0, 2.0, 2.0]$ for tanks 1, 2, 3 and 4 respectively. In addition to the challenge of NMP control, both 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 both control systems, with no controller-specific tuning or compensation.

Comparative evaluations are conducted across 10 simulation scenarios, with perturbations defined as follows:

- Coefficient γ_1 pertaining to the first pump degrades at a constant rate of $-0.001/s$.
- Coefficient γ_2, associated with the second pump undergoes a step-wise drop of 20% mid-simulation.

The 9th and 10th simulation, in addition to baseline scenario, feature a disturbance added at $t = 30s$; the campaign representing a stress-test of performance of both controllers under model mismatch and providing a validation of the robustness claim of LCA-RLS-KAN.

III. Controller Design & Stability Analysis

A. LTV-MPC Design

Linear Time-Varying MPC (LTV-MPC) uses the Jacobian matrix A derived from the system equations above and the input matrix B evaluated following Johansson [15]. The controller follows standard state-input weight scaling guidelines [2], [3]. Operator Splitting Quadratic Problem solver was selected for the LTV-MPC formulation [4]. The prediction horizon was set to $N = 30$, prediction time interval to $\Delta t_p = 0.1s$, $Q = \text{diag}([40, 40, 5, 5])$ and $R = \text{diag}([0.001, 0.001])$. Prediction horizon, cost weights and constraints remained identical across all runs, with no preemptive tuning performed; detailed configuration is featured in [17].

B. Controller Design

1) Structure of LCA-RLS-KAN: The control law of the LCA-RLS-KAN architecture is given as:

$$u(t) = B^{+}\left(\dot{x}_{ref} + \Phi_{KAN}^{T}\hat{\theta} - K_P e(t) + K_I \int e(t)dt + u_L\right) \quad (2)$$

where B^{+} is the left pseudoinverse of B, considering $B \in \mathbb{R}^{4x2}$ possesses full column rank in NMP operating point, $\dot{x}_{ref}$ is the feedforward term, Φ_{KAN} represents the nonlinear regressor vector, $\hat{\theta}(t)$ are the RLS-estimated linear weights. K_I represents the gain of integral action, K_P is the proportional gain and u_L represents Lipschitz continuous constraining the voltage, accounting for the error.

To demonstrate the transparency of the KAN architecture, the distilled control law for u_1 extracts into a single algebraic form, given by $u_1 = -0.097x_1 - 11.552x_2 + 0.24x_3 + ...$ The full law is reproduced in [17]; negative coefficients on x_1 and x_2 confirm physically expected damping, while quadratic terms encode the learned NMP coupling — auditability structurally unavailable in optimization-based controllers. The modifications ensuring boundedness and ISS properties are described in next subsection.

2) Lipschitz-Continuous RLS Update Law: To enable the adaptation of the controller, on top of the baseline KAN-distilled control law, three modifications were done.

Firstly, the parameter vector $\hat{\theta}$ is updated online using the continuous-time RLS law with forgetting factor λ [18], [19]:

$$\dot{\hat{\theta}} = \Gamma \Phi P e$$
$$\dot{\Gamma} = \rho\Gamma - \Gamma\Phi\Phi^{T}\Gamma$$

where:

$$\rho \approx \frac{1-\lambda}{T_a}$$

T_a representing the adaptation interval and $\tilde{\theta}$ representing estimated gains. The discrete implementation follows a standard RLS discretization (see supplementary material [17] for details).

While the original non-linear parameters learned by KAN had been frozen, the linear weights update via RLS with a forgetting factor, set to a positive constant ($0 < \lambda < 1$).

Secondly, integral action regulates the cumulative deviation from the standpoint. The error undergoes a low-pass filter to attenuate high-frequency components of the tracking error of the control response.

Finally, Lipschitz continuity as a function of time:

$$u_L = -K_L \frac{e(t)}{e_{max}} + slew_{lim}(-0.1, 0.1)$$

where $e(t) = x - x_{ref}$ and $slew_{lim}(-0.1, 0.1)$ denotes a slew rate limiter of 0.1 cm; thus, Lipschitz continuity was enforced based on deviation from the target state [20], constraining the controller's output, resulting in decreased actuator fatigue. As such, the control law remains restrained and tractable, verifying the claim of interpretability of the architecture made above.

3) Lyapunov Candidate: The energy state of the system is therefore governed by the adaptive laws outlined in literature [18], [19], [21], where the error propagates the adaptation and integral action engagement:

$$V(e, \tilde{\theta}, \tilde{d}) = \frac{1}{2}e^{T}Pe + \frac{1}{2}\tilde{\theta}^{T}\Gamma^{-1}\tilde{\theta} + \frac{1}{2}\tilde{d}^{T}Q_I\tilde{d} \quad (3)$$

where $e = x - x_{ref}$, $\tilde{d}$ denotes the integral action, $\tilde{\theta}$ is defined as parameter estimation error and Q_I is chosen such that $Q_I^{-1}P = K_I$. Since P, Γ^{-1}, and Q_I are assumed symmetric

positive definite (SPD), the Lyapunov candidate is positive definite and radially unbounded. Defining:

$$z = [e \ \tilde{\theta} \ \tilde{d}]^T$$

there exist positive constants $c_1, c_2 > 0$ such that:

$$c_1 ||z||^2 \leq V(z) \leq c_2 ||z||^2$$

Therefore, by Khalil Lemma 4.3 [20], V satisfies class-$\mathscr{K}_\infty$ bounds.

C. Stability Analysis

Taking the derivative of the Lyapunov candidate by trajectories of the system:

$$\dot{V} = e^T P \dot{e} + \tilde{\theta}^T \Gamma^{-1} \dot{\tilde{\theta}} + \tilde{d}^T Q_I \dot{\tilde{d}} + \frac{1}{2} \tilde{\theta}^T \dot{\Gamma^{-1}} \tilde{\theta}$$

Error dynamics are established as:

$$\dot{e} = \dot{x} - \dot{x}_{ref}$$
$$\dot{x} = f(x) + Bu$$

We define the plant dynamics $f(x)$ by KAN approximation as:

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

where $\varepsilon = d^* + \delta(x)$, d^* representing constant bias, $\delta(x)$ are time-varying perturbations and θ^* are ideal RLS gains. Then,

$$u = B^+ (\dot{x}_{ref} - Ke - \Phi^T \hat{\theta} - \hat{d})$$

where $\hat{\theta}$ and $\hat{d}$ are controller-approximated linear weights and integral action respectively. Assuming feasible reference trajectories, the projection residual $(BB^+ - I)(...)$ due to underactuation is bounded and absorbed into the lumped disturbance $\delta(x)$ [17]. Set $\theta^* - \hat{\theta} = -\tilde{\theta}$ and $d^* - \hat{d} = -\tilde{d}$, then:

$$\dot{e} = -\Phi^T \tilde{\theta} - \tilde{d} - Ke + \delta(x)$$

$K = K_P + K_L; K_L$ is the Lipschitz gain. Since ideal gains are assumed constant, $\dot{\tilde{\theta}}(t) = \dot{\hat{\theta}}(t)$. The gain update law is [18]:

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

The derivative of the integral action is:

$$\dot{\tilde{d}} = Q_I^{-1} P e = K_I e$$

In continuous-time, it is defined:

$$\rho \approx \frac{1 - \lambda}{T_a}$$

where T_a represents the adaptation interval. The continuous-time gain update law and its inverse therefore are:

$$\dot{\Gamma} = \rho \Gamma - \Gamma \Phi \Phi^T \Gamma$$
$$\dot{\Gamma^{-1}} = \Phi \Phi^T - \rho \Gamma^{-1}$$

Substituting the error dynamics and update laws yields:

$$\dot{V} = -e^T P (Ke - \delta) + \frac{1}{2} \tilde{\theta}^T \Phi \Phi^T \tilde{\theta} - \frac{1}{2} \rho \tilde{\theta}^T \Gamma^{-1} \tilde{\theta}$$

Using Cauchy-Schwarz inequality and Young's inequality, term $e^T P \delta$ is bounded:

$$e^T P \delta \leq ||e|| \, ||P|| \, ||\delta|| \leq \lambda_{max}(P) \, ||e|| \, ||\delta||$$
$$\lambda_{max}(P) \, ||e|| \, ||\delta|| \leq \frac{\mu}{2} ||e||^2 + \frac{\lambda_{max}(P)}{2\mu} ||\delta||^2$$

To ensure boundedness of the regressors, KAN terms were normalized. Hence, regressor terms can be expressed as:

$$\frac{1}{2} \tilde{\theta}^T \Phi \Phi^T \tilde{\theta} \leq \frac{1}{2} \Phi_{max}^2 \left|\left| \tilde{\theta} \right|\right|^2$$

where $\Phi_{max} = sup_{x \in \Omega} ||\Phi(x)||$, Ω representing a compact set of bounded tank heights, references and bounded integral state after slew rate limiting. The normalization is chosen such that $||\Phi(x)|| \leq \Phi_{max} \forall x \in \Omega$. Combining the terms:

$$-\frac{1}{2} \rho \Gamma^{-1} \left|\left| \tilde{\theta} \right|\right|^2 + \frac{1}{2} \Phi_{max}^2 \left|\left| \tilde{\theta} \right|\right|^2 = -\frac{1}{2} \left(\rho \Gamma^{-1} - \Phi_{max}^2 \right) \left|\left| \tilde{\theta} \right|\right|^2$$

To satisfy positive definiteness of the error term:

$$-e^T P K e + \frac{\mu}{2} ||e||^2 < 0$$
$$\lambda_{min}(PK) > \frac{\mu}{2}$$

Choose $\mu \in (0, \ 2\lambda_{min}(PK))$; for simplicity, set $\mu = \lambda_{min}(PK)$.

Due to exponential forgetting, $\Gamma(t)$ remains bounded on Ω and uniformly positive definite; therefore Γ^{-1} is bounded on Ω and $\lambda_{min}(\Gamma^{-1}) > 0$ (see [19]). It is sufficient:

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

which is ensured by the boundedness of the compact set; Φ_{max} is finite due to normalization, state bounds and slew rate limiting. Thus, the final form of the derivative is:

$$\begin{aligned} \dot{V} \leq & -\left(\lambda_{min}(PK) - \frac{\mu}{2} \right) ||e||^2 \\ & -\frac{1}{2} \left(\rho \lambda_{min}\left(\Gamma^{-1}\right) - \Phi_{max}^2 \right) \left|\left| \widetilde{\theta} \right|\right|^2 + \frac{\lambda_{max}(P)}{2\mu} ||\delta||^2 \end{aligned} \tag{4}$$

To prove ISS, set $\omega(t) := \delta(t)$. Composite functions are defined as:

$$||\zeta||^2 = \left|\left| \tilde{\theta} \right|\right|^2 + ||e||^2$$
$$||\omega||_\infty = ||\delta||_\infty$$
$$||\omega||_\infty^2 = ||\delta||_\infty^2$$

Additionally:

$$\alpha = min \left\{ \left(\lambda_{min}(PK) - \frac{\mu}{2} \right), \frac{1}{2} \left(\rho \lambda_{min}\left(\Gamma^{-1}\right) - \Phi_{max}^2 \right) \right\}$$
$$\gamma = \frac{\lambda_{max}(P)}{2\mu}$$

All matrices P, Q_I, Γ are assumed SPD and bounded [17]. The Unscented Kalman Filter estimation residual is absorbed into the L_∞-bounded disturbance term $\delta(x)$, supported by filter consistency in tuning (NIS ≈ 3.995) and the parameter bounds observed in Table II.

Theorem 1 (Stability of LCA-RLS-KAN)**.** *Consider the closed-loop system formed by the plant and the LCA-RLS-KAN controller. Let controller gains K, P, Q_I, Γ be symmetric positive definite. Assume the ideal weights are bounded ($||\theta^*|| \leq \theta_{max}$) and disturbances $\omega(t) := \delta(t)$ are bounded such that $\omega \in L_\infty$. If the damping coefficient satisfies $\rho\lambda_{min}(\Gamma^{-1}) > \Phi^2_{max}$, then the time derivative of the Lyapunov function satisfies:*

$$\dot{V} \leq -\alpha\,||\zeta||^2 + \eta\,||\omega||^2$$

with $\alpha \in \mathscr{K}_\infty$, $\eta \in \mathscr{K}$. By Khalil Definition 4.7 and Theorem 4.19 [20], the system is Input-to-State Stable.

Consequently, the tracking error $e(t)$ and parameter estimation error $\tilde{\theta}(t)$ are Globally Uniformly Ultimately Bounded, converging to the residual set characterized by the dissipation inequality:

$$\dot{V} \leq \frac{\lambda_{max}(P)}{2\mu}\,||\delta||^2$$

i.e., outside a ball whose radius depends on $||\delta||_\infty$, $\dot{V}$ is strictly negative.

Corollary 1 (Asymptotic Tracking)**.** *In the ideal case scenario where external disturbances vanish ($\delta \equiv 0$) and the KAN structure perfectly captures the plant dynamics ($\theta^* = 0$), the perturbation term becomes zero ($\omega = 0$). This yields $\dot{V} \leq -\alpha\,||\zeta||^2 \leq 0$. Invoking the LaSalle-Yoshizawa Theorem [22], it follows that:*

$$\lim_{t\to\infty} e(t) = 0$$

Corollary 2 (Robustness to Lack of Excitation)**.** *Under the conditions of Theorem 1, the tracking error $e(t)$ remains GUUB even in the absence of persistent excitation.*

Proof. The RLS update law with forgetting factor $\lambda < 1$ ensures that $\Gamma(t)$ remains bounded and uniformly positive definite by the exponential forgetting analysis of Theorem 1, which in turn guarantees that $\theta(t)$ remains bounded [18]. Thus, if $\tilde{\theta}$ does not converge to zero due to a lack of persistent excitation (PE), it remains in a bounded set. By Theorem 1, since the system is ISS with respect to external disturbances, tracking stability is preserved. □

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

IV. Results & Discussion

Simulations were conducted in Python 3.12.4 over a 50s horizon ($T_s = 3$ms, $N = 16666$). Runtime clock measured the complete control loop, including the estimator update and full control loop computation; divided by the number of discrete intervals, this yielded the computational results featured above. To ensure reproducibility, both controllers were evaluated on Acer Nitro 16 consumer laptop with Ryzen 7 7840HS CPU under identical multi-threaded execution environments. A joint Unscented Kalman Filter [23], [24], [25] was implemented via the FilterPy framework [26] to estimate the state vector $x_{UKF} = [h_{1\ldots4}, \gamma_{1,2}]$. The observer was tuned for nominal consistency (NIS ≈ 3.995) and tracking accuracy (RMSE$[h_{1\ldots4}] \approx 0.365$).

Comparative performance was validated via Monte-Carlo realizations (seed $= 42$) using identical zero-mean Gaussian noise realizations and parameter drift schedules. This protocol ensures that performance disparities are attributable to the control law rather than stochastic variability. Statistics are reported across eight baseline and two adversarial "stress-test" scenarios (Runs 9-10).

A. Baseline Scenario

Tables below summarize the first 8 Monte Carlo runs; disturbance injection runs (9-10) are analyzed separately. All statistics use median and interquartile range (IQR).

LCA-RLS-KAN incurs a deliberate tracking penalty on tanks 1–3, consistent with its design priority of actuator smoothness over strong setpoint pursuit; reduced IAE IQR across all tanks indicates greater deployment consistency.

Table I expands on the results reported, providing the median Integral Absolute Error (IAE) and IAE IQR on all 4 tanks across 8 runs on both controllers. It shows lower IAE for LTV-MPC on tanks 1-3, but reduced variability and lower tank 4 IAE for LCA-RLS-KAN.

TABLE I
Tracking Error Metrics (Lower is Better)

Controller	Tank	IAE Median $[cm \cdot step]$	IAE IQR $[cm \cdot step]$
LCA-RLS-KAN	1	50.269	**0.506**
LTV-MPC		**28.851**	0.806
LCA-RLS-KAN	2	40.936	**0.875**
LTV-MPC		**38.649**	1.459
LCA-RLS-KAN	3	37.788	**0.539**
LTV-MPC		**32.665**	1.151
LCA-RLS-KAN	4	**87.551**	**0.358**
LTV-MPC		116.496	1.924

Table II provided below illustrates the parameter estimation error. The parameters reviewed in table below are γ_1 and γ_2, corresponding to water flow ratios across pumps, which are not normally observable and rely on UKF estimation.

TABLE II
Parameter Estimation Error (Lower is Better)

Controller	Parameter	Median $\lvert e\rvert$	IQR $\lvert e\rvert$
LCA-RLS-KAN	γ_1	**0.044**	0.066
LTV-MPC		0.159	**0.021**
LCA-RLS-KAN	γ_2	**0.047**	0.064
LTV-MPC		0.140	**0.062**

Table II reveals a crucial part of the comparison: LCA-RLS-KAN achieves parameter estimation close to 3 times more precise compared to LTV-MPC, albeit with elevated IQR attributable to the forgetting factor λ, which by design discounts historical data to track non-stationary parameters — a known variance-bias trade-off in exponential forgetting RLS [18], consistent with Corollary 2. Moreover, LCA-RLS-KAN shows consistency across the estimation of both parameters, unlike LTV-MPC.

Table III shows the actuator effort demanded from both controllers to maintain their respective performance.

TABLE III
TOTAL VARIATION OF INPUT (LOWER IS BETTER)

Controller	Pump	Median (TV)	IQR (TV)
LCA-RLS-KAN	1	**1696.951**	**127.396**
LTV-MPC		9631.595	370.332
LCA-RLS-KAN	2	**1238.920**	288.913
LTV-MPC		4833.976	**247.705**

Table III displays LCA-RLS-KAN demonstrating substantially reduced control variation across both pumps, while maintaining comparable tracking performance.

Finally, runtime analysis revealed LCA-RLS-KAN achieving a median per-step execution of 22.512 ms (IQR: 0.374) compared to LTV-MPC's 113.424 ms (IQR: 5.246), representing a conservative 5x speedup.

Remark: As OSQP operates via C-based backend [4] while LCA-RLS-KAN runs exclusively on Python, the 5x runtime advantage is a conservative lower bound.

B. Adversarial Stress Test: Runs 9 and 10

The final two simulations introduce an additional step disturbance of 2.0 cm injected into Tank 2 at $t = 30s$ to rigorously probe robustness and cross-coupled setpoint tracking under pronounced modeling mismatch, emulating widely documented industrial disturbances such as valve obstruction, actuator degradation, and process noise. No controller re-tuning or structural modifications were performed, ensuring continuity of the initial simulation campaign.

The analysis focuses on tanks 2-4, directly and indirectly affected by the disturbance. Both controllers were evaluated under identical state and parameter settings, using the first of the four generated zero-mean Gaussian noise realizations to isolate performance differences attributable to controller design. The disturbance response of tank 2 is presented first.

As evidenced by Figure 1, under decaying plant conditions with additional disturbance, LCA-RLS-KAN exhibits faster convergence and reduced steady-state deviation compared to LTV-MPC. The most discriminating adversarial result is 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 [15], LCA-RLS-KAN maintains subsystem coherence, having started the restoration of h_3 at $t \approx 48s$ while LTV-MPC does not exhibit recovery of h_3 within the 50 second evaluation window.

Despite tank 4 being the upstream contributor to tank 2 inflow [15], LCA-RLS-KAN tracks tank 4 closer to its reference by achieving superior disturbance rejection on the pump 2 subsystem, indicating more coherent cross-coupled actuation (consistent with Table III's reduced TV).

Table IV provides the IAE and steady-state error (ESS) over the last 10 seconds for both control systems. Time-domain trajectories (Figs. 1, 2, 3) reveal trends not fully captured by ESS values. Thus, simulation campaigns support the claims of efficiency and robustness of the control law.

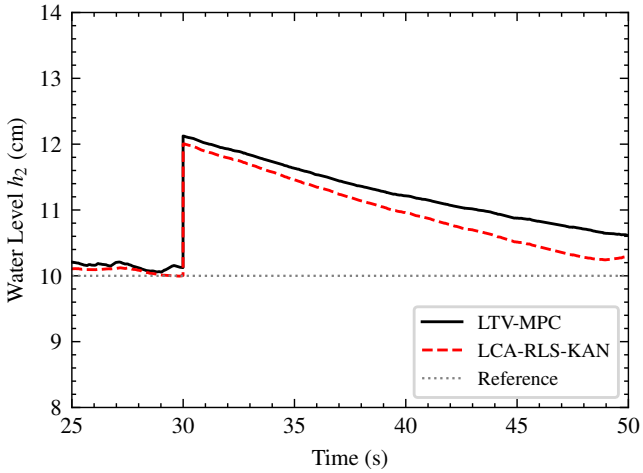

Fig. 1. **Tank 2 level tracking**. At $t = 30s$ a step disturbance of 2.0cm is applied, as shown in the plot. The disturbance rejection of two controllers on the affected tank itself is provided; LCA-RLS-KAN exhibits a steadier recovery compared to LTV-MPC.

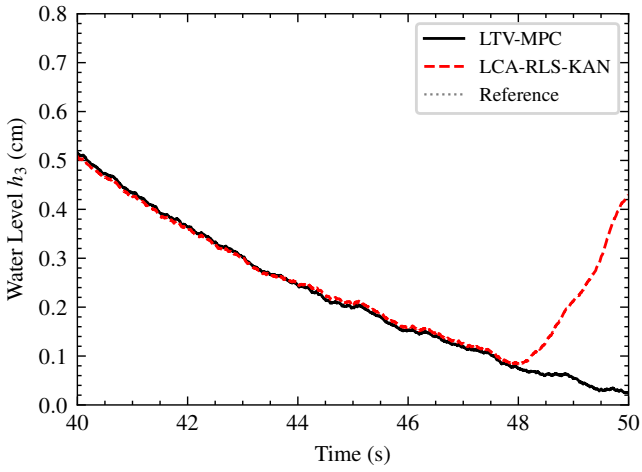

Fig. 2. **Tank 3 level tracking.** As tanks 2 and 3 are interconnected with a single pump, control of tank 3 represents a balancing challenge for both controllers: a necessity to reconcile the setpoint tracking of disturbed tank 2 with the draining of tank 3.

TABLE IV
TRACKING ERROR METRICS (LOWER IS BETTER)

Controller	Tank	IAE $[cm \cdot step]$	ESS [cm]
LCA-RLS-KAN	1	55.349	1.822
LTV-MPC		**31.906**	**0.560**
LCA-RLS-KAN	2	**58.342**	**0.344**
LTV-MPC		63.487	0.741
LCA-RLS-KAN	3	57.613	**1.823**
LTV-MPC		**57.049**	1.897
LCA-RLS-KAN	4	**94.386**	**3.351**
LTV-MPC		124.594	5.100

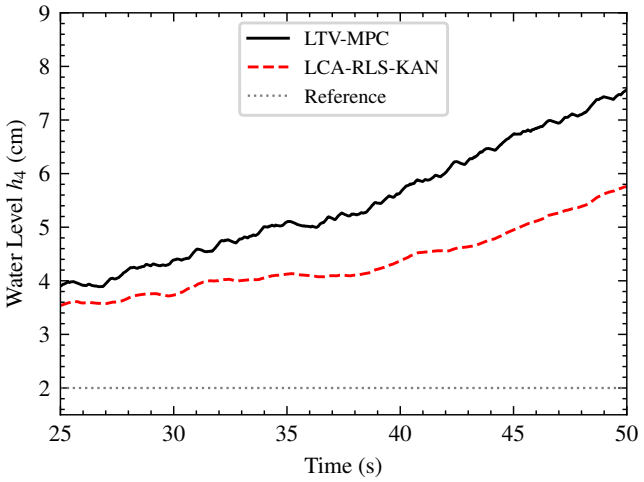

Fig. 3. **Tank 4 level tracking.** Reduced downward flow from tank 3 would increase the voltage needed to maintain the setpoint of tank 1, which in turn causes tank 4 level to rise. Therefore, lower level correlates with higher fidelity of multi-tank control and higher awareness of cross-coupling.

V. Conclusion

This paper presented a LCA-RLS-KAN architecture designed to bridge embedded feasibility and advanced nonlinear control. A Lyapunov-based analysis established ISS and GUUB, with asymptotic convergence under ideal conditions derived via LaSalle-Yoshizawa theorem and robustness under lack of PE proven. The theoretical guarantees explicitly account for bounded estimation error and parameter adaptation.

The simulation environment emulated documented industrial plant degradation scenarios. A Monte Carlo-style evaluation ensured identical stochastic noise realizations across controllers to enable controlled comparison.

In nominal testing, the proposed architecture achieved tracking performance comparable to LTV-MPC while reducing actuator variation (TV of input) by approximately a factor of five and runtime by a factor of five conservatively. Under adversarial disturbance conditions, corresponding to injected tank-level perturbations under decaying plant dynamics, LCA-RLS-KAN preserved subsystem coherence and restored cross-coupled state trajectories more effectively within the observed horizon. In particular, recovery dynamics of interconnected tanks remained bounded and exhibited upward convergence trends where LTV-MPC trajectories continued to degrade. These results indicate improved disturbance allocation management across coupled subsystems and more efficient actuation under plant degradation.

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 regressive control law suitable for constrained embedded systems.

References

[1] J. L. Jerez, P. J. Goulart, S. Richter, G. A. Constantinides, E. C. Kerrigan, and M. Morari, "Embedded online optimization for model predictive control at megahertz rates," *IEEE Transactions on Automatic Control*, vol. 59, no. 12, pp. 3238–3251, 2014.

[2] D. Q. Mayne, J. B. Rawlings, C. V. Rao, and P. O. Scokaert, "Constrained model predictive control: Stability and optimality," *Automatica*, vol. 36, no. 6, pp. 789–814, 2000.

[3] F. Borrelli, A. Bemporad, and M. Morari, *Predictive control for linear and hybrid systems*. Cambridge University Press, 2017.

[4] B. Stellato, G. Banjac, P. Goulart, A. Bemporad, and S. Boyd, "OSQP: An operator splitting solver for quadratic programs," *Mathematical Programming Computation*, vol. 12, no. 4, pp. 637–672, 2020.

[5] S. Boyd and L. Vandenberghe, *Convex optimization*. Cambridge University Press, 2004.

[6] S. Diamond and S. Boyd, "CVXPY: A Python-embedded modeling language for convex optimization," *Journal of Machine Learning Research*, vol. 17, no. 83, pp. 1–5, 2016.

[7] S. L. Brunton, J. L. Proctor, and J. N. Kutz, "Sparse identification of nonlinear dynamics with control (SINDYc)," *IFAC-PapersOnLine*, vol. 49, no. 18, pp. 710–715, 2016.

[8] H. Hose, A. Gräfe, and S. Trimpe, "Parameter-adaptive approximate MPC: Tuning neural-network controllers without retraining," in *6th Annual Learning for Dynamics & Control Conference*. PMLR, 2024, pp. 349–360.

[9] A. N. Kolmogorov, "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, 1957, pp. 953–956.

[10] V. I. Arnold, "On functions of three variables," in *Doklady Akademii Nauk*, vol. 114, no. 4. Russian Academy of Sciences, 1957, pp. 679–681.

[11] Z. Liu, Y. Wang, S. Vaidya, F. Ruehle, J. Halverson, M. Soljačić, T. Y. Hou, and M. Tegmark, "KAN: Kolmogorov-Arnold networks," *arXiv preprint arXiv:2404.19756*, 2024.

[12] G. Cruz, B. Renczes, M. C. Runacres, and J. Decuyper, "State-space Kolmogorov Arnold networks for interpretable nonlinear system identification," *IEEE Control Systems Letters*, 2025.

[13] T. Kim, A. Girard, and I. Kolmanovsky, "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*, ser. Proceedings of Machine Learning Research. PMLR, 04–06 Jun 2025, pp. 1115–1126.

[14] R. L. Plackett, "Some theorems in least squares," *Biometrika*, vol. 37, no. 1/2, pp. 149–157, 1950.

[15] K. H. Johansson, "The quadruple-tank process: A multivariable laboratory process with an adjustable zero," *IEEE Transactions on Control Systems Technology*, vol. 8, no. 3, pp. 456–465, 2000.

[16] S. Skogestad and I. Postlethwaite, *Multivariable feedback control: analysis and design*. John Wiley & Sons, 2005.

[17] A. Salkimbayev, "Supplementary materials for "Robust Adaptive Kolmogorov-Arnold Neural Control"," *Zenodo preprint https://doi.org/10.5281/zenodo.19578563*, 2026.

[18] P. A. Ioannou and J. Sun, *Robust adaptive control*. PTR Prentice-Hall Upper Saddle River, NJ, 1996, vol. 1.

[19] J.-J. E. Slotine, W. Li *et al.*, *Applied nonlinear control*. Prentice-Hall Englewood Cliffs, NJ, 1991, vol. 199, no. 1.

[20] H. K. Khalil and J. W. Grizzle, *Nonlinear systems*, 3rd ed. Prentice-Hall Upper Saddle River, NJ, 2002.

[21] K. S. Narendra and A. M. Annaswamy, *Stable adaptive systems*. Courier Corporation, 2012.

[22] M. Krstic, P. V. Kokotovic, and I. Kanellakopoulos, *Nonlinear and adaptive control design*. John Wiley & Sons, Inc., 1995.

[23] R. E. Kalman, "A new approach to linear filtering and prediction problems," *Transactions of the ASME–Journal of Basic Engineering*, vol. 82, no. Series D, pp. 35–45, 1960.

[24] E. A. Wan and R. Van Der Merwe, "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)*. IEEE, 2000, pp. 153–158.

[25] S. J. Julier and J. K. Uhlmann, "Unscented filtering and nonlinear estimation," *Proceedings of the IEEE*, vol. 92, no. 3, pp. 401–422, 2004.

[26] R. R. Labbe, *Kalman and Bayesian Filters in Python*. Chapman & Hall/CRC, 2014. [Online]. Available: https://github.com/rlabbe/Kalman-and-Bayesian-Filters-in-Python