POLITECNICO MILANO 1863

SCUOLA DI INGEGNERIA INDUSTRIALE
E DELL'INFORMAZIONE

Executive Summary of the Thesis

Optimal Control Strategies for the Quadruple Tank Plant: from Linear Quadratic Regulator to Adaptive Scenario-based Model Predictive Control

Laurea Magistrale in Automation and Control Engineering – Ingegneria dell'Automazione

Author: Feruzakhon Sulaymonova

Advisor: Prof. Fredy Orlando Ruiz Palacios

Co-advisors: Dr. Andres Felipe Cordoba Pacheco, Prof. Eva Masero Rubio

Academic year: 2025–2026

1. Introduction

Multivariable control problems characterized by strong loop interactions, actuator saturation and parameter uncertainty require strategies that handle constraints directly. To address this the thesis investigates five advanced control architectures using a laboratory quadruple-tank process. This benchmark provides genuine multivariable coupling and enforces realistic hardware limitations: sensor noise, actuator saturation, and outflow valve coefficients that drift between $-25\,\%$ and $-10\,\%$ of their nominal values due to fouling, creating a persistent parametric mismatch that fixed models cannot compensate [2]. Starting from an identified grey-box nonlinear model the study designs and experimentally validates a Linear Quadratic Regulator with integral action (LQI), nominal Model Predictive Control (MPC), MPC with successive linearization (SL-MPC), scenario-based MPC (SC-MPC) and adaptive scenario-based MPC.

The primary objective is to establish a unified methodology anchored to the LQI baseline to quantify the performance trade-offs of these strategies. Robustness to parametric mismatch is addressed through scenario-based prediction and further improved by the adaptive architecture which uses online set-membership identification with zonotopic bounding to track time-varying dynamics and reduce control conservatism.

2. System Description and Modeling

The experimental plant is a Quanser quadruple water tank system. It consists of four interconnected tanks equipped with differential pressure sensors and discharge valves alongside two pumps driven by voltage commands. The adjustable flow split valves enable precise control over the multivariable transmission zeros.

For pump 1 a fraction γ_1 of the flow is directed to tank 3 while the remaining fraction $1 - \gamma_1$ goes to tank 2. Pump 2 splits its flow between tank 1 and tank 4. Tanks 1 and 3 drain into tanks 2 and 4 respectively. The experiments use the minimum-phase configuration ($\gamma_1 = \gamma_2 = 0.6$), which places the transmission zeros

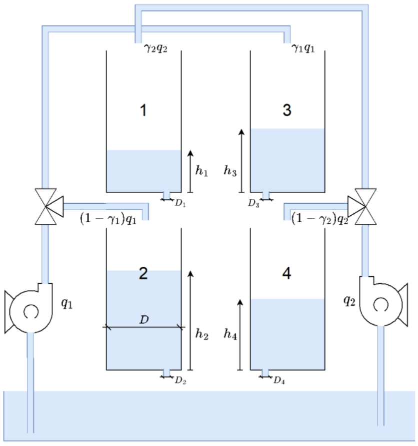

Figure 1: QTS schematic.

in the left half-plane and ensures the velocity-form MPC framework remains internally stable. The controlled outputs are the bottom tanks $y = [h_2 \; h_4]^\top$ and the manipulated inputs are the pump voltages $u = [V_1 \; V_2]^\top$. The nonlinear model is identified via a grey-box procedure that fits the four Torricelli discharge coefficients a_i and the pump flow maps $k_{1,2}$ to step-response data, yielding the identified valve intervals that motivate the uncertainty set in Section 3.

3. Control Architectures

Five controllers of increasing capability are designed under a unified velocity-form MPC framework to guarantee offset-free tracking without an explicit disturbance observer. All share identical Bryson-scaled cost weights, horizons $N = 60$, $N_c = 10$, and hardware constraints $u \in [4, 18]\,\mathrm{V}$, $|\Delta u| \leq 0.5\,\mathrm{V}$.

LQI. The baseline computes a state-feedback gain offline via the discrete algebraic Riccati equation on an augmented system with integrated output error. Voltage limits are enforced by post-hoc projection with anti-windup clamping of the integral state during saturation.

Nominal MPC. A condensed QP problem is solved online at each sample. The prediction model is fixed at the nominal equilibrium linearisation, which degrades during large transients due to the nonlinear Torricelli discharge dynamics.

SL-MPC. The Jacobians are recomputed at each sample using the current measured state [1], extending prediction validity without sacrificing QP problem convexity. The outflow derivatives

$$q_{\mathrm{out},i}(t) = A_i \sqrt{2g\, h_i(t)}, \tag{1}$$

are updated online.

SC-MPC. Outlet valve coefficients drift due to fouling; each is modelled as lying in

$$[0.75\, A_i^{\mathrm{d}},\, 0.90\, A_i^{\mathrm{d}}] \tag{2}$$

This interval, spanning $-25\,\%$ to $-10\,\%$ of the design value, was determined from repeated grey-box identification runs on the physical plant across different fouling states. The sample count $K = 39$ follows from the Schildbach bound [4]

$$K \geq \frac{\rho}{\epsilon} - 1 \tag{3}$$

at risk level $\varepsilon = 5\%$ with $n_u = 2$ input channels. Worst-case deviations across scenarios tighten output constraints; a slack variable preserves feasibility.

Adaptive SC-MPC. Although the fixed interval (2) guarantees robustness, it is deliberately conservative: if the true valve coefficients lie near the upper end of the interval the scenario spread unnecessarily tightens constraints and degrades tracking — in the symmetric setpoint test this conservatism alone costs 23 s of additional settling time on the second tank relative to LQI. The adaptive architecture removes this conservatism by learning the true uncertainty set online. Online set-membership identification [3] maintains an axis-aligned box Θ_k updated via Minkowski expansion ($\lambda_{\mathrm{exp}} = 0.005$, applied at every 60 steps) and strip intersection with noise bound $\varepsilon_{\mathrm{smi}} = 0.35\,\mathrm{cm}$. A 40-sample guard period inhibits learning during startup; a 100-sample cooldown after reference changes prevents biased estimates from transient finite-difference derivatives; and a per-tank derivative threshold ($|\dot{h}_i| < 0.3\,\mathrm{cm/s}$) provides an additional guard against residual dynamics outside the cooldown window. As Θ_k contracts, scenario spread narrows and the QP problem recovers performance lost to conservatism.

This design was not straightforward to deploy: in preliminary hardware tests without the cooldown or derivative threshold, the finite-difference derivatives during reference transients were corrupted by sensor noise, producing biased valve coefficient estimates that collapsed Θ_k to a degenerate interval and rendered the scenario-tightened QP problem infeasible. The guard period, cooldown and derivative threshold were introduced iteratively to suppress learning during these transient windows, ensuring the box remains well-conditioned for scenario generation.

4. Experimental Evaluation

The five controllers form a deliberate progression: each architecture is motivated by a limitation of its predecessor observed on the real plant. LQI establishes the speed baseline but cannot enforce constraints predictively. Nominal MPC adds constraint handling but its fixed linearisation degrades during large fills because Torricelli dynamics are strongly nonlinear far from the operating point. SL-MPC corrects this by updating the Jacobian online but still uses a nominal valve model. SC-MPC addresses valve uncertainty explicitly but is conservative because it assumes the worst-case interval at all times. Adaptive SC-MPC removes that conservatism by learning the true interval from data. The tests below expose exactly these failure modes and quantify the gains at each step.

4.1. Symmetric Setpoint Tracking

Test 1 establishes the baseline comparison. LQI's lack of predictive constraint handling is immediately visible: pump 1 saturates at 18 V, producing the fastest settling ($\leq$ 26 s) but the highest energy (16 568 V^2 s) and largest input variation (71.4 V). The MPC controllers trade settling speed for smoother, constraint-respecting actuation.

Table 1: Performance metrics for Test 1 ($h_2 = h_4 =$ 10 cm)

Metric	Tank	LQI	Nom. MPC	SL-MPC	SC-MPC
IAE (cm s)	h_2	148.4	184.8	173.3	174.3
	h_4	121.3	117.8	114.7	117.9
RMSE (cm)	h_2	2.910	3.204	3.053	2.893
	h_4	2.542	2.421	2.383	2.345
Settling (s)	h_2	25.8	35.5	38.0	49.0
	h_4	24.6	25.8	24.4	30.3
Overshoot (%)	h_2	6.85	5.85	5.47	6.10
	h_4	6.94	3.67	3.94	4.43
Energy (V^2 s)	–	16 568	15 507	15 245	14 954
Total Var. (V)	–	71.4	64.8	56.2	62.7

LQI achieves the fastest settling ($\leq$ 26 s on both tanks) and the lowest IAE on h_2, but requires the highest control energy (16 568 V^2 s) and the largest input variation among the open-loop startup configurations, with pump 1 saturating at 18 V immediately after controller engagement. Nominal MPC reduces energy by avoiding aggressive actuation but is the slowest on h_2 (35.5 s), reflecting the penalty on input increments. SL-MPC achieves the lowest total variation (56.2 V) and the lowest IAE on h_4 (114.7 cm s), confirming that online re-linearisation produces smoother trajectories without sacrificing steady-state accuracy. SC-MPC records the lowest control energy (14 954 V^2 s) but the most conservative h_2 settling time (49.0 s) due to constraint tightening from the fixed uncertainty interval.

4.2. Step and Pulse Reference Tracking

The step and pulse tests expose the limitation of Nominal MPC's fixed linearisation. On the large $0 \to 10$ cm fill, where the system is furthest from the nominal operating point, Nominal MPC is the slowest controller (39.8 s on h_2 in Test 2). SL-MPC closes this gap by updating the Jacobian online, achieving the smoothest actuation (76.1 V total variation in Test 3) without sacrificing IAE on h_4 (55.8 cm s, best among all controllers).

Table 2: Global metrics for Test 2 ($0 \to 10 \to 12 \to$ 10 cm, sequential).

Metric	Tank	LQI	Nom. MPC	SL-MPC	SC-MPC
IAE (cm s)	h_2	77.0	102.7	106.6	90.2
	h_4	56.7	57.3	55.8	79.0
RMSE (cm)	h_2	0.876	1.120	1.128	0.886
	h_4	0.548	0.521	0.518	0.742
Energy (V^2 s)	–	31 281	29 747	29 842	29 118
Total Var. (V)	–	126.8	131.5	125.5	144.8

Table 3: Global metrics for Test 3 ($0 \to 10 \to 12 \to$ 10 cm, simultaneous).

Metric	Tank	LQI	Nom. MPC	SL-MPC	SC-MPC
IAE (cm s)	h_2	72.7	89.9	95.1	84.0
	h_4	48.9	48.3	46.7	77.6
RMSE (cm)	h_2	1.223	1.422	1.583	1.167
	h_4	0.698	0.688	0.672	1.115
Energy (V^2 s)	–	19 648	18 159	18 165	17 607
Total Var. (V)	–	88.5	80.5	76.1	84.8

Across both tests LQI achieves the lowest combined IAE and tightest tracking on h_2, with SL-MPC consistently delivering the smoothest actuation (lowest total variation in Test 3: 76.1 V) and the best h_4 IAE in Test 2 (55.8 cm s). SC-MPC achieves the lowest control energy in both tests but produces the largest h_4 RMSE under simultaneous excitation (1.115 cm), reflecting sensitivity to coupling when both channels demand flow simultaneously. Nominal MPC is the slowest on the initial $0 \to 10$ cm fill (39.8 s in Test 2) because its fixed linearisation degrades furthest from the nominal operating point.

4.3. Differential Reference Tracking and Feasibility

Test 4 stresses the controllers with conflicting reference commands, exposing the fundamental limitation of LQI: without predictive constraint awareness it drives Tank 1 to overflow under differential excitation and must be excluded. The three MPC controllers remain safe but all exhibit elevated steady-state error on h_4 ($\approx$ 0.34–0.36 cm), revealing that cross-coupling under opposing flow demands acts as an unmeasured disturbance that no fixed linear model fully compensates. Results for the three MPC controllers are shown in Table 4.

Table 4: Global metrics for Test 4 (differential references).

Metric	Tank	Nom. MPC	SL-MPC	SC-MPC
IAE (cm s)	h_2	102.2	98.2	95.1
	h_4	56.2	54.4	72.5
RMSE (cm)	h_2	1.079	1.092	0.985
	h_4	0.530	0.528	0.700
SS error (cm)	h_4	0.341	0.361	0.361
Energy (V^2 s)	–	27 636	27 473	26 766
Total Var. (V)	–	127.8	126.4	136.0

All three MPC controllers exhibit an elevated steady-state error on h_4 ($\approx$ 0.34–0.36 cm) compared to symmetric tests, consistent with cross-coupling acting as an unmeasured disturbance when the two channels request opposing flow changes. SC-MPC achieves the lowest h_2 IAE (95.1 cm s) while SL-MPC provides the smoothest actuation (lowest total variation: 126.4 V).

This test also motivated a systematic feasibility analysis of the admissible reference space. Attempting $h_{2,\mathrm{ref}} = 9$ cm, $h_{4,\mathrm{ref}} = 12$ cm requires $\bar{h}_1 \approx 30$ cm, which exceeds the 27 cm physical overflow limit; the experiment confirmed continuous accumulation in Tank 1 until the safety interlock triggered termination. The reversed configuration ($h_{2,\mathrm{ref}} = 12$ cm, $h_{4,\mathrm{ref}} = 9$ cm) is fully feasible. This establishes $h_{2,\mathrm{ref}} \geq h_{4,\mathrm{ref}}$ as a necessary condition for steady-state operability, arising from the physical flow routing and independent of the controller architecture. Five validated asymmetric operating points within the feasible region $9 \leq h_{\mathrm{ref}} \leq 15$ cm are documented in the thesis. The admissible reference space is mapped in Figure 2.

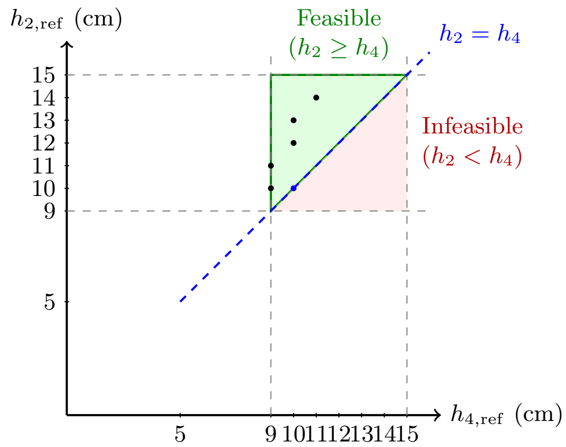

Figure 2: Feasibility map of the admissible reference space.

4.4. Disturbance Rejection

The disturbance rejection test reveals the one scenario where SC-MPC's conservatism is an advantage. Tighter constraint margins keep the optimiser alert to boundary proximity, producing the fastest local recovery from both injections (10.3 s and 10.7 s) despite SC-MPC having the worst settling time in setpoint tracking. Segmented recovery metrics are reported in Table 5.

Table 5: Disturbance response metrics. Peak = maximum deviation from setpoint; T_r = recovery time.

Metric	LQI	SC-MPC	SL-MPC	Nom. MPC
h_2 injection ($t = 100$ s)				
IAE	12.77	11.75	15.20	12.46
RMSE	0.540	0.554	0.721	0.554
Peak (cm)	2.031	2.000	2.957	1.905
T_r (s)	15.8	10.3	13.2	11.8
h_4 injection ($t = 150$ s)				
IAE	11.70	8.62	10.60	9.38
RMSE	0.482	0.428	0.462	0.437
Peak (cm)	1.674	1.715	1.805	1.776
T_r (s)	15.4	10.7	10.2	11.5

LQI recovers with the lowest global IAE (90.7 and 55.4 cm s on h_2 and h_4 respectively over the full experiment window) but the longest local recovery times (15.8 s and 15.4 s), reflecting its aggressive yet predictively unguided response.

SC-MPC achieves the fastest local recovery from both injections (10.3 s and 10.7 s) and the lowest per-event IAE, confirming that robust optimisation is advantageous when unmeasured disturbances push the system toward constraint boundaries. SL-MPC records the largest peak deviation on the h_2 injection (2.957 cm) because the Jacobian update amplifies the initial transient before stabilising. When the injection shifts the operating point abruptly, the updated Jacobian briefly reflects the disturbed rather than the recovering dynamics, causing a one-to-two sample under-correction before the linearisation tracks the true state. Nominal MPC delivers the smallest h_2 peak (1.905 cm) with balanced recovery. The apparent contradiction arises because LQI's aggressive steady-state actuation suppresses background drift effectively, keeping global IAE low, while its lack of predictive recovery planning produces slower return after the impulsive disturbance.

4.5. Adaptive SC-MPC vs SC-MPC

Having established that SC-MPC's fixed uncertainty interval (2) causes unnecessary conservatism when true valve coefficients lie near the upper end of the range, the adaptive controller is tested under identical conditions to isolate the learning benefit. During the first 40 samples the SMI guard period is active and both controllers behave identically, applying the same conservative constraint tightening derived from the initial uncertainty box $\Theta_0 = [0.75\,A_i^{\mathrm{d}},\, 0.90\,A_i^{\mathrm{d}}]$. Once the guard period expires, the adaptive controller begins contracting Θ_k from steady-state measurements. Table 6 quantifies the resulting improvement.

Table 6: SC-MPC vs Adaptive SC-MPC: benefit of online set-membership learning.

Metric	SC-MPC	Adaptive MPC	Improvement
IAE h_2 (cm s)	129.4	110.5	15 %
IAE h_4 (cm s)	153.6	149.8	3 %
RMSE h_2 (cm)	2.332	1.979	15 %
RMSE h_4 (cm)	2.880	2.856	1 %
Overshoot h_2 (%)	6.98	5.13	26 %
Overshoot h_4 (%)	5.43	4.64	15 %
Settling h_2 (s)	37.1	38.4	–
Settling h_4 (s)	35.4	30.0	15 %
Energy (V^2 s)	15 523	14 955	4 %
Total Var. (V)	73.6	67.8	8 %

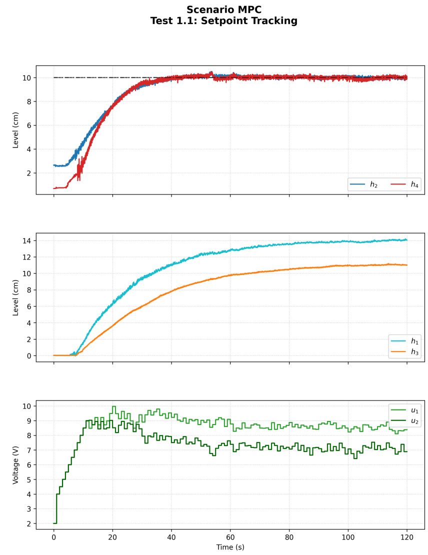

Figure 3: Setpoint tracking on the real plant (SC-MPC).

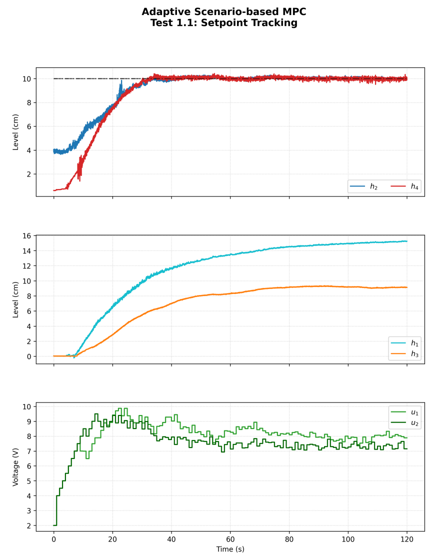

Figure 4: Setpoint tracking on the real plant (Adaptive SC-MPC).

As Θ_k contracts, the scenario spread narrows and the constraint tightening margins decrease, granting the QP problem optimizer additional freedom to minimise the tracking cost. This manifests as a 15 % reduction in IAE and RMSE on h_2, a 26 % reduction in overshoot on h_2, and a 4 % saving in control energy. The settling time on h_2 is marginally slower for the adaptive controller (38.4 s vs 37.1 s), which falls within experimental repeatability and is not considered significant. The improvement is concentrated on h_2 because Tank 2 is the primary output of Pump 1, whose valve coefficient exhibits the larger drift relative to the initial uncertainty interval. Learning on h_4 is limited by the noisier

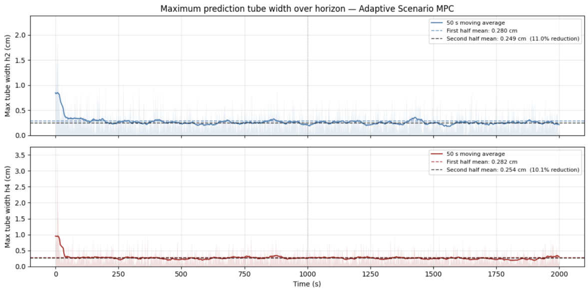

Figure 5: Scenario envelope.

Tank 4 sensor, whose signal-to-noise ratio during steady-state is insufficient to drive meaningful strip intersection within the 120 s experiment window, which explains the 3 % improvement on h_4 versus 15 % on h_2.

The convergence of the online learning is quantified by the maximum scenario envelope width (see Figure 5) over the $N = 60$ step prediction horizon, evaluated in a 2000 s nonlinear grey-box simulation with additive measurement noise reproducing the closed-loop behaviour of the physical plant. The envelope contracts from an initial peak of approximately 0.9 cm to a steady-state value of approximately 0.25 cm within the first 60 s — a 70 % reduction driven by the first valid SMI observations after the guard period expires. Under continuous steady-state operation, the spread continues to contract slowly (10–11 % further reduction over 2000 s); when pulse reference changes are applied, the cooldown suspends learning during each transition, limiting the residual contraction to 5–9 %. The steady-state spread of 0.25 cm is small relative to the operating range and constraint margins, confirming that the learned uncertainty set supports effective constraint tightening without excessive conservatism.

5. Conclusions

The experimental campaign establishes a clear performance hierarchy across five controllers validated on real hardware. LQI excels at transient speed but lacks constraint prediction; nominal MPC improves energy efficiency at the cost of transient accuracy; SL-MPC delivers the smoothest actuation through online re-linearisation; SC-MPC provides the strongest disturbance rejection near constraint boundaries; and adaptive SC-MPC combines online learning with scenario optimisation to progressively recover from conservatism, achieving a 15 % IAE reduction and 26 % overshoot reduction over fixed SC-MPC. The structural feasibility constraint $h_{2,\text{ref}} \geq h_{4,\text{ref}}$ was identified and experimentally confirmed: violating it causes Tank 1 overflow regardless of the controller used, and must be enforced at the supervisory level, as it is intrinsic to the plant flow routing.

Future directions include replacing the grey-box mass balance regressor with a black-box ARX model to eliminate dependence on offline plant knowledge, extending the simulation validation of the adaptive SC-MPC to the full test suite (multi-step, differential and disturbance scenarios), improving the Tank 4 sensor calibration to accelerate parameter box contraction, adopting a full zonotopic uncertainty representation to capture cross-parameter correlations, and incorporating active excitation during steady-state periods to reduce the dependence on reference changes as the primary source of identification data.

References

[1] Julian Berberich, Johannes Köhler, Matthias Müller, and Frank Allgöwer. Linear tracking MPC for nonlinear systems part i: The model-based case. *IEEE Transactions on Automatic Control*, 67(9):4426–4441, 2022.

[2] Karl Henrik Johansson. The quadruple tank process: A multivariable laboratory process with an adjustable zero. *IEEE Transactions on Control Systems Technology*, 8(3):456–465, 2000.

[3] Esteban Meza, Alessandro Del Duca, Fredy Ruiz, and Carlos Ocampo-Martinez. Adaptive scenario-based predictive control: A set membership learning approach. *IFAC-PapersOnLine*, 59(26):323–328, 2025.

[4] Georg Schildbach, Lorenzo Fagiano, Christian Frei, and Manfred Morari. The scenario approach for stochastic model predictive control with bounds on closed-loop constraint violations. *Automatica*, 50(12):3009–3018, 2014.