POLITECNICO
MILANO 1863

SCUOLA DI INGEGNERIA INDUSTRIALE
E DELL'INFORMAZIONE

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

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

Author: **Feruzakhon Sulaymonova**

Student ID: 243084
Advisor: Prof. Fredy Orlando Ruiz Palacios
Co-advisors: Dr. Andres Felipe Cordoba Pacheco, Prof. Eva Masero Rubio

Academic Year: 2025-26

Abstract

This thesis investigates advanced control strategies for multivariable constrained systems using a laboratory quadruple-tank process as a benchmark. Five control approaches of increasing complexity are developed, implemented and compared: 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 with set-membership identification. All controllers are tested experimentally and in simulation via MATLAB/Simulink.

Starting from an identified grey-box model, this work proposes a unified methodology for controller design, tuning and validation. First, the LQI is suitably tuned using a common tunning strategy to provide a reference baseline. Second, MPC is formulated in Velocity-form to ensure offset-free tracking and to handle water level and actuator rate constraints. Moreover, successive linearization approach is used to extend the validity of the prediction model beyond a single operating point. Third, robustness to parametric uncertainty is addressed through a scenario-based approach and further enhanced by implementing adaptive scenario-based MPC, which combines online set-membership identification with zonotopic parameter bounding to handle time-varying system parameters. Experimental evaluation of controllers includes reference tracking, disturbance rejection and robustness tests under varying operating conditions. Results show that LQI achieves excellent nominal performance with minimal computational cost, while advanced MPC strategies provide improved performance due to constraint handling and robustness at the expense of higher complexity. Adaptive scenario-based MPC demonstrates particular effectiveness in maintaining performance under parametric model variations while reducing the conservatism of SC-MPC (due to fixed uncertainty bounds) through online refinement of the zonotopic parameter uncertainty set.

Keywords: model predictive control, quadruple-tank plant, successive linearization, scenario-based MPC, set-membership identification, zonotopes

Abstract in lingua italiana

Questa tesi analizza strategie di controllo avanzate per sistemi multivariabili soggetti a vincoli, utilizzando come benchmark un processo di laboratorio a quattro serbatoi. Vengono sviluppati, implementati e confrontati cinque approcci di controllo con complessità crescente: regolatore lineare quadratico con azione integrale (LQI), Model Predictive Control nominale (MPC), MPC con linearizzazione successiva (SL-MPC), MPC basato su scenari (SC-MPC) e MPC basato su scenari adattivo con identificazione set-membership. Tutti i controllori sono testati sperimentalmente e in simulazione tramite MATLAB/Simulink.

A partire da un modello grey-box identificato, il lavoro propone una metodologia unificata per la progettazione, la taratura e la validazione dei controllori. In primo luogo, l'LQI viene opportunamente tarato mediante una strategia di tuning comune, così da fornire un riferimento di base. In secondo luogo, l'MPC è formulato in forma di velocità per garantire un inseguimento offset-free e gestire i vincoli sui livelli d'acqua e sui ratei degli attuatori. Inoltre, la linearizzazione successiva viene adottata per estendere la validità del modello di predizione oltre un singolo punto di lavoro. In terzo luogo, la robustezza rispetto all'incertezza parametrica è affrontata tramite un approccio basato su scenari e ulteriormente migliorata implementando l'SC-MPC adattivo, che combina identificazione set-membership online e vincolamento zonotopico dei parametri per gestire parametri di sistema variabili nel tempo.

La valutazione sperimentale dei controllori include prove di inseguimento del riferimento, reiezione dei disturbi e test di robustezza in condizioni operative variabili. I risultati mostrano che l'LQI raggiunge prestazioni nominali eccellenti con un costo computazionale minimo, mentre le strategie MPC più avanzate offrono prestazioni migliori grazie alla gestione dei vincoli e alla robustezza, a fronte di una maggiore complessità. L'SC-MPC adattivo si dimostra particolarmente efficace nel mantenere le prestazioni in presenza di variazioni parametriche del modello e nel ridurre il conservatorismo dell'SC-MPC (dovuto a limiti di incertezza fissi) tramite il raffinamento online dell'insieme di incertezza zonotopico dei parametri.

Parole chiave: controllo predittivo, impianto a quattro serbatoi, linearizzazione succes-

siva, MPC su scenari, identificazione set-membership, zonotopi

Contents

List of Figures

List of Tables

1 | Introduction

1.1. Motivation and Problem Statement

Multivariable control problems with strong loop interactions, hard actuator limits and uncertain parameters are still a recurring source of performance loss and commissioning effort in industrial automation [1]. In many plants, a single manipulated variable affects multiple regulated quantities through shared material or energy paths. Therefore, independent SISO tuning cannot guarantee good behaviour when operating conditions change or when constraints become active [2]. When a SISO approach is applied to a coupled multivariable plant, the resulting controller either becomes too conservative to exploit the available actuator authority, leading to slow setpoint transitions and reduced throughput, or too aggressive, causing constraint violations that trigger safety actions and increase actuator wear and energy consumption [2]. For these reasons, industrial control design increasingly focuses on methods that treat coupling and constraints directly rather than compensating for them after the fact. In practice, this also means making controller selection on the basis of measurable metrics, such as tracking error, energy use and actuator activity, since these quantities determine product variability, cycle time, operating cost and maintenance burden [3].

For model-based controllers, the practical difficulty is that models may not always be accurate since physical parameters such as valve coefficients and flow gains vary with operating conditions, drift over time due to wear and fouling, and differ from their nominal values whenever the plant moves away from the identification point. Even a carefully identified model becomes less accurate as the operating point changes, since linearization-based predictors are only valid locally and nonlinear effects grow with the distance from the design equilibrium. Sensors introduce additional bias and noise, and unmeasured disturbances further shift the effective operating conditions. The consequence is that inaccurate predictions lead to suboptimal transients, performance deterioration and constraint violations that would not appear in simulation, motivating controllers that explicitly account for parametric uncertainty rather than treating the model as exact [1].Even when a controller performs well in nominal conditions, small model mismatches can lead to steady-state offset or actuator saturation that degrades reliability. Offset-free operation is therefore not only a theoretical property but a requirement for long runs, since persistent modelling errors and disturbances can be rejected without manual retuning [4], [5]. In this context, velocity-form MPC provides a practical route to offset-free tracking that fits well with constrained optimization and real-time implementation [4], [5]. This combination is attractive in practice because it supports a straightforward commissioning workflow: constraints are enforced explicitly, offset is removed structurally, and controller can be similarly tuned to account for performance and wear requirements.

To address these challenges in a controlled environment, this work employs the quadruple-tank process as a benchmark. Its two pumps influence the controlled levels through adjustable flow splits and internal cascade paths, which create genuine multivariable coupling and can yield nonminimum-phase behaviour depending on the valve configuration [6], [7]. At the same time, the plant enforces realistic limitations: pump voltages saturate, measurement signals contain noise, and the effective outflow parameters vary with operating conditions and component ageing. This makes the benchmark suitable for analysing not only tracking performance but also constraint handling, sensitivity to uncertainty and the impact of modelling choices on closed-loop behaviour. Consequently, we judge the results by whether each strategy respects the constraints without being overly cautious, and whether it keeps the control actions within acceptable wear and energy limits.

Within this benchmark, several MPC variants offer different practical benefits. First, nominal MPC can handle constraints explicitly but relies on a prediction model linearized around an operating point that becomes less accurate as the operating point moves. Successive linearisation extends the validity of a linear MPC formulation by updating the prediction linear model online, which can improve performance when nonlinearities are relevant. Second, SC-MPC introduces a structured way to enforce constraints under bounded parametric uncertainty providing a trade-off between performance and robustness [8].Third, Adaptive SC-MPC can go a step further by updating parameter bounds online through set-membership identification, aiming to preserve constraint satisfaction while reducing conservatism when the data supports it [5]. From an process control perspective, these three variants correspond to different risk and cost profiles, ranging from simpler to more robust formulations.

This work investigates how the MPC design choices translate into measurable differences of performance for managing a real multivariable plant. The comparison of controllers is centred on offset-free velocity formulations and is conducted under conditions that matter in practice: actuator saturation, measurement noise, operating-point changes and parametric uncertainty. Therefore, trade-offs between tracking quality, constraint handling and control effort serves as guidelines to select the controller needed [2], [4], [8]. In this sense, the contribution of this work is not limited to benchmarking, since the same guideline can be reused for choosing and tuning a controller for industrial MIMO plants where commissioning time, safety margins and actuator lifetime are key constraints [9].

1.2. Thesis Organization

This section outlines the thesis structure.

Chapter 2, *Literature Review*, surveys multivariable process control and the quadruple-tank benchmark [7]. It reviews the evolution of Model Predictive Control from its industrial roots [2] to scenario-based and adaptive scenario-based formulations [8], [10], with particular attention to set membership identification [11] and zonotopic set representations [12] used for robust adaptation.

Chapter 3, *System Description and Experimental Setup*, describes the laboratory plant, its modelling assumptions and the experimental workflow. It presents the grey-box identification used to estimate the plant parameters and introduces the velocity-form formulation adopted to achieve offset-free tracking [4], [5] without explicit disturbance estimation.

Chapter 4, *Optimal Control Strategies*, presents the controller design and tuning methodology for the full set of control approaches considered in this thesis. It is organised into four parts: (i) the LQI controller is used as baseline and the common tuning procedure used to support a consistent comparison, (ii) the deterministic linear MPC in velocity form, which covers the nominal MPC formulation and the successive linearisation strategy [13], [14], (iii) the SC-MPC which introduces a stochastic MPC formulation to handle bounded parametric uncertainty [8] and (iv) the adaptive SC-MPC with set-membership identification and zonotopic parameter uncertainty updates that combines predictive control of model parameters and online learning to reduce conservatism as information accumulates [12], [15].

Chapter 5, *Discussion of Results*, reports a comparative analysis of controllers' evaluation. Performance is assessed in both the linearised and nonlinear models as well as the physical plant, using key indicators that quantify tracking error, control effort and robustness to modelling parametric uncertainties.

Chapter 6, *Conclusions and Future Work*, summarises the main trade-offs between computational complexity and closed-loop performance. It discusses practical implications for industrial deployment and outlines directions for future research, including learning-based modelling and control extensions [16], [17].

2 | Literature Review

This chapter presents a state-of-the-art review in multivariable process control and builds the theoretical background required for the controllers developed in this thesis. The review begins with LQI, which provides the optimal unconstrained baseline against which all subsequent controllers are compared. It then outlines the development of MPC from its early industrial formulations to recent stochastic and adaptive variants. Particular attention is given to offset-free tracking formulations, to successive linearization as a strategy for extending linear MPC to nonlinear operating regimes, and to set membership identification as a practical tool for representing and refining parametric uncertainty in adaptive SC-MPC.

2.1. Quadruple Tank System as a Benchmark

The Quadruple Tank System (QTS) process was originally introduced by K. H. Johansson [7]. The system consists of four interconnected water tanks with two pumps and adjustable valve configurations (see schematic in Figure 2.1) that enable precise control over the multivariable transmission zeros. The adjustable zero property is particularly significant: by modifying the valve split ratios, the system can be configured to have zeros in either the left or right half-plane. For minimum-phase configurations (where the sum of valve ratios exceeds unity), aggressive control is feasible. For nonminimum-phase configurations, fundamental bandwidth limitations arise due to right-half-plane zeros, requiring careful controller tuning to maintain stability. This thesis employs a minimum-phase configuration to focus on constraint handling and uncertainty management rather than bandwidth limitations. This unique feature allows researchers to configure the system for either minimum-phase or nonminimum-phase behavior, making it particularly valuable for testing controller robustness and performance limits. Moreover, the plant is typically equipped with differential pressure sensors that output analog voltages, and these signals can exhibit gain mismatch, bias, and slow drift due to ambient conditions [18].

Therefore, the quadruple tank exhibits several characteristics that mirror industrial multivariable processes: strong coupling between control loops, actuator saturation constraints, measurable output disturbances, and nonlinear dynamics due to varying outflow rates. These features have made it a standard testbed in control laboratories worldwide, with numerous studies demonstrating various control architectures ranging from classical decentralized PID control to advanced robust and adaptive scenario-based MPC formulations [6], [19].

Beyond its educational value, the system has been utilized to illustrate various control architectures including Internal Model Control (IMC), Dynamic Matrix Control (DMC) and multivariable robust control [20]. Research has demonstrated its utility in comparing centralized and decentralized schemes; for instance, a Distributed Model Predictive Control (DMPC) framework was proposed in [21], showing that communication between local controllers approaches centralized performance. Furthermore, the study [22] have compared traditional linear MPC with Neural Network-based variants (NN-MPC) to manage the trade-off between tracking accuracy and computational complexity. Other authors have explored sliding mode control to handle parametric uncertainty in valve coefficients [23] and H_∞ loop shaping for robust rejection of external disturbances [24].

Figure 2.1: Lab-scale quadruple tank system in experimental setup.

2.2. Linear Quadratic Regulator

While many studies on the quadruple tank system utilize decentralized PI or PID controllers as baseline [6], [7], [25], [26], [27], while via Linear Quadratic Regulator (LQR) one can draw fundamental conclusions regarding the behavior of infinite and finite control horizon strategies. LQR employs the optimal state-feedback solution for an infinite-horizon quadratic objective. By utilizing the same weighting matrices Q and R across all controllers, a fair comparison can be done allowing for a rigorous analysis of how the transition from a closed-form offline solution to a constrained online optimization affects the transient and steady-state responses.

The LQR control law for linear, discrete time-invariant systems is obtained by minimizing a quadratic performance index that balances regulation accuracy against input effort [1]. Such that

$$x_{k+1} = A_d x_k + B_d u_k \tag{2.1}$$

$$J = \sum_{k=0}^{\infty} \left(x_k^\top Q x_k + u_k^\top R u_k \right) \tag{2.2}$$

where the feedback control law is $u_k = -K x_k$ with gain K obtained from solving the discrete algebraic Riccati equation. Under the usual stabilizability and detectability conditions, the Riccati solution yields a stabilizing gain and provides robustness properties that make LQR a useful baseline in multivariable regulation problems [1]. The same controller design extends naturally to multivariable plants by choosing Q and R to reflect the relative importance of regulation errors and actuation levels across channels.

Pure LQR is a regulation design around the chosen equilibrium and it does not guarantee zero steady-state tracking error when constant disturbances or model mismatch are present. A common extension is LQR controller augmented with integral action (LQI), where integral action is introduced by augmenting the state with the integrated output error. With the measured output

$$y_k = C x_k \tag{2.3}$$

the integral state can be defined as

$$x_{I,k+1} = x_{I,k} + (r_k - y_k) \tag{2.4}$$

so that offset errors between the desired reference r_k and the output y_k accumulate until the feedback compensates them. This augmentation embeds the internal model needed for reference tracking and disturbance rejection, enabling offset-free behaviour for step commands under persistent biases [1]. In experimental settings, this property is essential because even small parameter errors, sensor bias or unmodelled losses can lead to a persistent steady-state offset. In compact form, the LQI design is obtained by applying the same LQR procedure to the augmented model

$$\tilde{x}_k = \begin{bmatrix} x_k \\ x_{I,k} \end{bmatrix}, \qquad \tilde{x}_{k+1} = \underbrace{\begin{bmatrix} A_d & 0 \\ -C & I \end{bmatrix}}_{\tilde{A}} \tilde{x}_k + \underbrace{\begin{bmatrix} B_d \\ 0 \end{bmatrix}}_{\tilde{B}} u_k + \underbrace{\begin{bmatrix} 0 \\ I \end{bmatrix}}_{\tilde{E}} r_k \tag{2.5}$$

which leads to the final control law

$$u_k = -\tilde{K}\tilde{x}_k \tag{2.6}$$

with $\tilde{K}$ given by the Riccati solution for the augmented dynamics [1].

A practical challenge in LQR and LQI design is the selection of the weighting matrices Q and R. Bryson type scaling provides a systematic starting point by normalizing states and inputs with respect to acceptable deviations, after which weights can be refined to match the desired trade off between speed and actuation level [1]. In practice, this refinement is often empirical rather than derived from a closed form rule, because the weight selection can reflect hardware limits, wear considerations and the specific operating regime of the plant. For LQI designs, the integral states are commonly weighted more softly to avoid overly aggressive transients, while additional scaling factors can be introduced to penalize rapid voltage variations beyond what simple normalization would suggest, improving actuator longevity. Empirical adjustment is also typical when the controller is later embedded in learning or identification loops, where the link between weights and the evolving model is not explicit and the final choice is guided by measured closed-loop behaviour. However, the resulting feedback is unconstrained. When actuator limits become active, the closed-loop behaviour may deviate from the nominal design and integral action can introduce windup if saturation persists, so saturation and windup must be handled outside the optimal design, motivating constrained formulations such as MPC [1].

Therefore, MPC can then be seen as a constrained extension of the same idea, replacing the infinite horizon objective with a finite horizon optimization solved online and enforcing explicit bounds on inputs, states or outputs [2]. This connection is useful for a comparative study because it clarifies what changes when moving from LQI to MPC: the tuning objective remains broadly similar, but feasibility and constraint management become part of the MPC controller design rather than an afterthought.

2.3. Model Predictive Control

MPC has become the reference technology in process industries since its industrial adoption in the 1970s because it treats multivariable coupling and constraints directly within the control design [3]. From a historical perspective, MPC did not emerge as a single unified development. Rather, it consolidated ideas that matured along several lines of work: state-space optimal control, generalized predictive control in the input-output tradition, and industrial process control practices where constraint handling and operational objectives were central in practice. This consolidation was also supported by links between dynamic programming (Hamilton-Jacobi-Bellman) and the maximum principle, which clarified how constrained optimal control laws can be approximated through repeated finite-horizon optimization. In this sense, modern MPC can be read as a computationally viable proxy for infinite-horizon optimal control, enriched with explicit constraints and a receding-horizon implementation.

More specifically, at each sampling instant, MPC computes the current input by solving a finite horizon optimal control problem (FHOCP) based on a prediction model. Only the first control move is applied, then the optimization is repeated at the next sample in a receding horizon fashion. This mechanism allows the controller to trade tracking performance against actuation effort while ensuring feasibility with respect to explicit bounds on inputs, states and outputs, which is a core practical advantage over unconstrained optimal feedback laws [2].

A nominal linear MPC formulation starts from a discrete-time linear model

$$x_{k+1} = A_d x_k + B_d u_k, \qquad y_k = C x_k \tag{2.7}$$

and solves at time k an optimization over a prediction horizon N of the form

$$\min_{\{u_{k+i|k}\}_{i=0}^{N-1}} \sum_{i=0}^{N-1} \left(\left\| y_{k+i|k} - r_{k+i} \right\|_Q^2 + \left\| u_{k+i|k} \right\|_R^2 \right) + \left\| x_{k+N|k} \right\|_P^2 \tag{2.8}$$

subject to the prediction dynamics and constraints such as

$$u_{\min} \le u_{k+i|k} \le u_{\max}, \qquad y_{\min} \le y_{k+i|k} \le y_{\max}, \qquad x_{\min} \le x_{k+i|k} \le x_{\max}, \qquad i = 0, \ldots, N-1 \tag{2.9}$$

In this nominal setting, the controller assumes a fixed prediction model and does not explicitly account for uncertainty, so constraint satisfaction and performance are only as reliable as the prediction model in the relevant operating region. A finite-horizon optimal solution does not, by itself, provide a closed-loop guarantee. Two properties are typically singled out in MPC theory: recursive feasibility and convergence or stability. Recursive feasibility means that if the optimization problem is feasible at the initial time then feasibility can be maintained at all subsequent sampling instants under the receding-horizon policy. Convergence or stability concerns the behaviour of the closed-loop state, typically requiring that trajectories approach the desired equilibrium. These properties are not automatic consequences of solving (2.8) subject to (2.9). They are usually enforced through terminal ingredients, namely a terminal cost $V_f(\cdot)$ and a terminal constraint set $\mathcal{X}_f$, chosen so that $\mathcal{X}_f$ is positively invariant under a suitable local control law and the resulting value function admits a Lyapunov-type decrease argument .

MPC typically requires an offset-free mechanism to achieve zero steady-state error under persistent plant-model mismatch. While the grey-box model captures the dominant physics, it remains an approximation of the real plant's complex dynamics. Parameter uncertainty, nonlinearities, unmodelled losses, sensor bias and operating point dependent effects can therefore lead to a steady-state offset if the controller relies only on the nominal prediction model [5].

Offset-free tracking is a primary requirement for the quadruple tank benchmark, as permanent errors often arise from Torricelli's Law nonlinearities and valve fouling. In [5], Pannocchia provides a coherent framework for classifying offset-free formulations into three categories: error integration, disturbance models with observers and velocity-form representations. Classical error integration, as used in LQI, augments the state with integrated output error but often requires separate anti-windup logic. Disturbance-based approaches such as [28], implemented for the quadruple tank, augment the model with integrating states representing virtual disturbances estimated via an observer. This thesis adopts the MPC problem in velocity-form, which utilizes control increments Δu as decision variables. This approach implicitly embeds an integral action without introducing an explicit disturbance state and observer, which simplifies the offset-free layer [5]. More in detail, with

$$\Delta u_k = u_k - u_{k-1} \tag{2.10}$$

the prediction model in velocity-form can be written as follows

$$\bar{x}_{k+1} = \begin{bmatrix} A_d & B_d \\ 0 & I \end{bmatrix} \bar{x}_k + \begin{bmatrix} B_d \\ I \end{bmatrix} \Delta u_k, \qquad y_k = \begin{bmatrix} C & 0 \end{bmatrix} \bar{x}_k, \qquad \bar{x}_k = \begin{bmatrix} x_k \\ u_{k-1} \end{bmatrix}. \tag{2.11}$$

A practical consequence is that setpoint changes can be handled without explicitly recomputing steady-state targets for (x, u) at every reference update. In incremental form, where the decision variable is the control increment $\Delta u_k = u_k - u_{k-1}$ rather than the absolute input u_k, stage cost, the per-step contribution to the objective function

$$\ell(\Delta y_{k+i|k},\, \Delta u_{k+i|k}) = \left\|\Delta y_{k+i|k}\right\|_Q^2 + \left\|\Delta u_{k+i|k}\right\|_R^2, \tag{2.12}$$

naturally vanishes when the output error is driven to zero and the increments converge to zero, which matches the steady-state regime sought in tracking problems [4], [5]. In practice, it also provides a convenient handle to penalize actuator activity directly via Δu, which is often more aligned with wear and smoothness requirements than penalizing u alone [5].

The velocity-form MPC problem solved at each sampling instant is

$$\min_{\{\Delta u_{k+i|k}\}_{i=0}^{N-1}} \sum_{i=0}^{N-1} \left\|\Delta y_{k+i|k}\right\|_Q^2 + \left\|\Delta u_{k+i|k}\right\|_R^2 \tag{2.13}$$

subject to input increment and output constraints, where $\Delta u_k = u_k - u_{k-1}$ is the control increment and $\Delta y_{k+i|k}$ is the predicted output deviation from the reference.

Nominal linear MPC performs well when the process operates close to the linearisation point, but its prediction quality can degrade when nonlinearities become relevant or when the operating point moves significantly. To manage the nonlinearities inherent in the discharge dynamics, this work implements MPC with successive linearization (SL-MPC) [13], [14]. Unlike full nonlinear MPC (NMPC) which involves non-convex optimization, successive linearization updates the prediction model at each sampling step using the Jacobian of the nonlinear vector field evaluated at the current measured state [14]. Theoretically, this allows the system to track the steady-state manifold toward the reachable equilibrium, a property that has been related to exponential stability under uniform controllability assumptions on the linearized pairs. In practice, stability is verified through simulation and experimental testing under the specific operating conditions of the plant. [14]. This technique provides a significant computational advantage by reducing the online task to a convex Quadratic Program (QP) that can be solved in milliseconds, making it more suitable for real-time hardware implementation than iterative nonlinear MPC algorithms [29]in which non-convex optimization problem is solved directly over the nonlinear model at each sampling instant and requires iterative solvers with no guarantee of convergence within the available sampling time. For a nonlinear discrete-time model

$$x_{k+1} = f(x_k, u_k) \tag{2.14}$$

the system is linearized at each sample around the current state x_k and the previously applied input u_{k-1}. Defining deviation variables

$$\delta x_k = x_k - \bar{x}_k, \qquad \delta u_k = u_k - \bar{u}_k, \qquad \delta y_k = y_k - \bar{y}_k, \tag{2.15}$$

with linearization point $(\bar{x}_k, \bar{u}_k) = (x_k, u_{k-1})$, the local Jacobians

$$A_k = \left.\frac{\partial f}{\partial x}\right|_{(\bar{x}_k,\bar{u}_k)}, \qquad B_k = \left.\frac{\partial f}{\partial u}\right|_{(\bar{x}_k,\bar{u}_k)} \tag{2.16}$$

define the linearized predictor in deviation coordinates

$$\delta x_{k+1} \approx A_k\,\delta x_k + B_k\,\delta u_k + d_k, \tag{2.17}$$

where $d_k = f(\bar{x}_k, \bar{u}_k) - A_k\bar{x}_k - B_k\bar{u}_k$ is the affine correction that makes the linearized model exact at the current operating point. The MPC problem is then built using (A_k, B_k, d_k) for prediction over the horizon. A common practical choice is to keep the model frozen over the prediction horizon and update it at the next sample, which limits complexity while capturing the dominant operating-point dependence [13], [14].

The benefit of successive linearization is that it preserves the QP structure of linear MPC while updating the prediction model to reflect the current operating point, which targets operating-point-dependent mismatch. This is especially useful when local gains vary with the state, since a fixed linear model may lose accuracy during large transients. The price is practical rather than conceptual: each sample requires computing Jacobians and making design choices on horizon freezing (keeping the linearization point fixed over the prediction horizon rather than updating it at each step), refresh frequency (how often the Jacobian is recomputed relative to the sampling time), and treatment of the affine correction (whether the constant offset term from the Taylor expansion is retained or discarded in the QP). Compared with the fixed-model case, stability and robustness require additional assumptions and are typically supported by tuning and validation under constraints and measurement. In this thesis, the nominal prediction model, which is linearised once around the operating point, serves as the baseline, while successive linearization extends this to the model around other operating points, introduced as an incremental modification that mitigates operating-point-induced mismatch without altering the underlying optimization problem.

2.4. Scenario-Based MPC

All mathematical models are essentially uncertain and real-world industrial control loops are inevitably affected by parametric uncertainty, unmodeled dynamics and disturbances. While MPC in velocity form ensures offset-free tracking by implicitly satisfying the Internal Model Principle for constant errors, it does not eliminate the impact of model mismatch during the prediction horizon. Inaccurate predictions can result in suboptimal transients, performance deterioration or constraint violations not observed in simulation.

To address these issues, Robust MPC methodologies, such as min-max and tube-based approaches [2], aim to guarantee stability and constraint satisfaction for all potential uncertain trajectories. However, satisfying constraints for every possible uncertainty value is often overly conservative, particularly when worst-case realizations are rare or when the uncertainty bounds are large [8]. Furthermore, the computational demands of these robust MPC methods often do not scale well with increasing model complexity, which can be limiting for real-time implementations [3].

Scenario-based MPC provides a more tractable compromise by enforcing constraints on a finite set of K sampled uncertainty realizations. This randomized optimization approach relaxes the deterministic worst-case requirement into a quantified risk of constraint violation, trading off strict satisfaction for improved cost function performance. The theoretical foundation of this method ensures that if a sufficient number of scenarios is considered, constraint satisfaction holds with a specific confidence level. By utilizing support-rank arguments, the design targets a prescribed expected violation probability that is explicitly related to the number of scenarios and the number of system input channels [8].

A convenient description starts from an uncertain prediction model parameterized by $\theta \in \Theta$,

$$x_{k+1} = f(x_k, u_k, \theta), \qquad y_k = h(x_k, \theta), \qquad \theta \in \Theta \tag{2.18}$$

where Θ collects the admissible parameters. In scenario-based MPC, a set of samples $\{\theta^{(j)}\}_{j=1}^{K}$ is drawn from a distribution supported on Θ and the constraints are enforced for all sampled models simultaneously [8]. The number of scenarios K is an explicit design choice and it controls the trade off between robustness and computational burden [3], [8].

For a prediction horizon N, each scenario generates its own predicted trajectory,

$$x^{(j)}_{k+i+1|k} = f\left(x^{(j)}_{k+i|k}, u_{k+i|k}, \theta^{(j)}\right), \qquad i = 0, \ldots, N-1, \quad j = 1, \ldots, K \tag{2.19}$$

while the input sequence $\{u_{k+i|k}\}$ is shared across scenarios, reflecting the fact that one control move must work for all considered realizations. The scenario-based MPC problem can be formulated as:

$$\min_{\{u_{k+i|k}\}} \sum_{i=0}^{N-1} \left(\left\| y_{k+i|k} - r_{k+i} \right\|_Q^2 + \left\| u_{k+i|k} \right\|_R^2 \right) + \left\| x_{k+N|k} \right\|_P^2 \tag{2.20}$$

$$\text{s.t. } x_{k|k}^{(j)} = x_k, \qquad j = 1, \dots, K \tag{2.21}$$

$$x_{k+i+1|k}^{(j)} = f\left(x_{k+i|k}^{(j)}, u_{k+i|k}, \theta^{(j)} \right), \qquad i = 0, \dots, N-1, \quad j = 1, \dots, K \tag{2.22}$$

$$u_{\min} \leq u_{k+i|k} \leq u_{\max}, \qquad i = 0, \dots, N-1 \tag{2.23}$$

$$y_{\min} \leq y_{k+i|k}^{(j)} \leq y_{\max}, \qquad i = 1, \dots, N, \quad j = 1, \dots, K \tag{2.24}$$

where $y_{k+i|k}^{(j)} = h(x_{k+i|k}^{(j)}, \theta^{(j)})$. Compared with the nominal MPC problem (2.8)-(2.9), the only structural change is that the constraints are replicated across scenarios which increases the problem size. A useful way to interpret the scenario approach is through the trade-off between conservatism and quantified risk. Robust MPC targets worst-case guarantees over an uncertainty set, which can lead to conservative inputs when the uncertainty description is wide. Scenario-based MPC replaces the worst-case requirement with a finite set of sampled constraints, leading to solutions that are typically less conservative while admitting explicit bounds on the probability of constraint violation [8]. In other words, the price paid for improved economic performance is not an unstructured loss of guarantees, but a controlled and user-tunable risk level expressed through the scenario sample size and confidence parameters [8]. This is sometimes described as a commuter's dilemma: accept worst-case conservatism or accept a quantified probability of violation in exchange for improved performance.

The selection of the number of scenarios K, often referred to as *sample complexity*, is a critical design choice that governs the trade-off between computational burden and probabilistic safety. In the literature, two primary frameworks exist for determining K based on the desired nature of the guarantee: the *a priori* bound of Campi and Garatti, which fixes K before data collection based on the risk level ϵ and confidence β, and the *a posteriori* approach of Calafiore, which tightens the bound after solving the scenario program [1], [30], [31].

The first approach, derived from the theory of Scenario Programming, provides a confidence-based guarantee for the feasibility of the entire predicted trajectory in a single optimization step. This method ensures that the probability of constraint violation V_t remains below a risk level ϵ with a high reliability level $1 - \beta$, where β is a small confidence parameter (e.g., 10^{-9}). According to the results established in [1], an explicit upper bound for K is given by:

$$K \geq \frac{2}{1-\epsilon} \log(\beta^{-1} + mN + 1) \tag{2.25}$$

where m represents the number of system inputs and N is the prediction horizon. While this explicit bound is highly useful for theoretical certification, it is often considered too conservative for real-time control because it ignores the multi-stage structure of the MPC problem and the benefits of receding horizon implementation [8]. To reduce this conservatism while maintaining the (ϵ, β) guarantee, a tighter, implicit condition is often used in practice:

$$\sum_{j=0}^{d-1} \binom{K}{j} \epsilon^j (1-\epsilon)^{K-j} \leq \beta \tag{2.26}$$

where d is the number of degrees of freedom involved in the constraints, typically associated with the number of optimization variables [8], [30].

The second approach, specifically tailored for real-time closed-loop operation, focuses on controlling the average rate of constraint violations over an infinite time horizon [32], [33]. By exploiting the fact that the sampling procedure is repeated at each time step, a significantly lower sample complexity can be achieved. Under this framework, the required number of scenarios is computed as [8]:

$$K \geq \frac{\rho}{\epsilon} - 1 \tag{2.27}$$

In this expression, $\rho \leq n_u$ represents the support rank of the problem, which is defined as the dimension of decision space minus the dimension of the unconstrained subspace of the constraint [33]. Since the same K scenarios are enforced at every receding-horizon step, the expected constraint violation probability satisfies

$$\mathbb{E}\left[V_t \mid y_t\right] \leq \frac{\rho}{K+1}, \tag{2.28}$$

where $V_t = \mathbb{P}\left[\theta \notin \Theta^* \mid y_t\right]$ is the probability that a new unseen uncertainty realisation violates the constraints at time t

Scenario-based MPC applies the same sampling based construction at each step, either by resampling scenarios online or by refreshing a subset of them, so that uncertainty handling remains aligned with the current operating conditions [8]. This is particularly relevant when uncertainty is dominated by parameter variations that affect constraint margins, since the controller must remain feasible while maintaining a low rate of violations. Scenario-based MPC has been demonstrated successfully in application such as energy management and microgrids, where uncertainty on demand and renewable generation is naturally described in probabilistic terms and constraint violations must remain rare [34]. Within the scope of this thesis, the scenario-based formulation is introduced as an intermediate step between nominal MPC and adaptive SC-MPC architecture, since it provides a structured way to improve constraint robustness under bounded parametric uncertainty.

2.5. Adaptive scenario-based MPC with Set Membership Identification

While scenario-based MPC handles fixed parametric uncertainty robustly, many real systems exhibit time-varying parameters due to aging, fouling, seasonal variations or changing operating conditions. Adaptive scenario-based MPC addresses this challenge by combining online identification with predictive control so that the prediction model is updated as new data becomes available, preserving feasibility and performance under model changes [10]. In this sense, identification can be seen as a continuous calibration of a model-based digital twin used to predict the constrained evolution over the horizon [10]. Set Membership identification (SMI) fits naturally in this context when uncertainty is described through credible bounds rather than distributional assumptions, since it maintains a set of models consistent with the measurements and the assumed error bounds [11].

SMI differs from classical stochastic estimation, for instance recursive least squares or Kalman filtering, because it does not commit to a single parameter estimate with an associated covariance [11]. Instead, it characterizes all parameter vectors that can explain the observed data under a bounded error description. This is particularly relevant when disturbances are bounded but their statistical properties are unknown, drifting or difficult to justify. For control design, the main benefit is that uncertainty is quantified directly in a form that can be propagated to predictions and embedded into constraint handling. As measurements accumulate, the feasible set tends to contract when the data are informative and consistent, while it may expand when excitation is poor or when the underlying dynamics change. This behaviour provides an adaptive notion of model confidence that is well aligned with robust predictive control.

$$y_{t+1} = \theta^\top z_t + d_{t+1}, \quad z_t = [y_t, \ldots, y_{t-n_a+1}, u_t, \ldots, u_{t-n_b+1}]^\top \tag{2.29}$$

where y_{t+1} is the plant output, z_t is the regressor vector, and d_{t+1} represents the modelling error. Here n_a denotes the number of past output lags and n_b the number of past input lags in the ARX regressor

$$\|d_t\|_\infty \leq \epsilon_d \tag{2.30}$$

The modeling error is assumed to be unknown-but-bounded [35], which defines the following family of models consistent with the data through the chosen bound ϵ_d:

$$S_t = \left\{\theta : \left|y_t - \theta^\top z_{t-1}\right| \leq \epsilon_d\right\}, \quad FSS_t = FSS_{t-1} \cap S_t \tag{2.31}$$

where the Feasible System Set (FSS) is updated recursively as the intersection of the strips S_t generated by the measurements.

While the SM framework above is formulated for general ARX-type models, it applies equally to any regression structure where the output depends linearly on unknown parameters. For adaptive SC-MPC, this property is exploited by casting the quadruple tank mass balance equations as a linear regression in the outlet valve coefficients, so that the zonotopic SM machinery can be applied directly to the physical model parameters.

A zonotope is defined by its center p and generator matrix H, offering a convenient compromise between representational tightness and computational tractability [12], [15].

$$\mathcal{Z} = p \oplus HB^m = \{p + H\xi : \|\xi\|_\infty \leq 1\} \tag{2.32}$$

The FHOCP enforces the constraints across K scenarios obtained by sampling parameters from the current uncertainty set $\mathcal{Z}_t$ [8], [33], [35].

$$\min_{U_N} \sum_{i=0}^{N-1} \|y^{(0)}_{i|t} - r_{t+i}\|^2_Q + \|u_{i|t}\|^2_R \tag{2.33}$$

$$\text{s.t.} \quad y^{(k)}_{i+1|t} = \left(\theta^{(t)}_k\right)^\top z^{(k)}_{i|t}, \quad y^{(k)}_{i|t} \in \mathbb{Y}, \quad u_{i|t} \in \mathbb{U}, \quad \forall\, k \in \{1, \ldots, K\} \tag{2.34}$$

$$\theta^{(k)}_t \sim \mathcal{Z}_t, \quad k = 1, \ldots, K \tag{2.35}$$

where $\Delta U_N = [\Delta u^\top_{0|t}, \ldots, \Delta u^\top_{N-1|t}]^\top$ is the sequence of control increments optimised over the horizon, and the K parameter scenarios $\{\theta^{(k)}_t\}^K_{k=1}$ are sampled from the current uncertainty set $\mathcal{Z}_t$. The prediction in 2.34 uses the noise-free regression; the modelling error d_{t+1} satisfying $\|d_t\|_\infty \leq \epsilon_d$ is accounted for implicitly through the scenario spread rather than included explicitly in the prediction constraint.

This structure supports the operations needed in SM identification, in particular linear mappings, Minkowski sums and strip intersections, with predictable computational cost [12], [15]. The complexity of a zonotope is often related to its order, defined as m/n_z, where m is the number of columns of the generator matrix H and n_z is the dimension of the parameter space. In adaptive scenario-based MPC implementations, order reduction is therefore used to keep the update and sampling steps compatible with real-time requirements. Recent studies report that low-order representations can already deliver effective learning and control performance in practice [36], [37]. In this thesis a first-order zonotope ($m = n_z = 4$) is used, where each generator corresponds to one valve coefficient, keeping the order reduction step unnecessary.

The combination of zonotopic SM identification with scenario-based MPC yields an adaptive predictive control architecture where uncertainty quantification and constraint handling are coupled by construction [33], [38]. At each time step, the parameter set is updated using the new measurement, typically contracting under consistent informative data and expanding when a change is detected. The scenario set used in the MPC problem then follows the current uncertainty level, which reduces unnecessary conservatism when the model is confident and increases protection when uncertainty grows. This mechanism avoids reliance on convergence to a single parameter estimate and is compatible with operating regimes where persistent excitation cannot be enforced. Applications on battery energy storage systems and benchmark studies report improved performance compared to non-adaptive strategies when parameters vary over time [35], [36].

Recent developments in learning-based control have explored the integration of data-driven models with MPC frameworks [16]. Recurrent neural networks (RNNs) have shown promise for learning complex dynamics from data, with stability conditions that support their inclusion within predictive control loops [39]. Physics-informed neural networks (PINNs) combine first-principles structure with data-driven learning, offering an intermediate modelling approach between purely mechanistic and purely black-box representations [17].

From a safety perspective, SMI provides a rigorous baseline for learning in control because it maintains explicit uncertainty bounds that can be propagated to predictions and constraints. This connection between SMI and safe learning remains an active research direction, especially for systems operating under safety-critical constraints.

3 | System Description and Experimental Setup

This chapter describes the laboratory quadruple tank process, the models used for controller design and the real-time implementation workflow.

3.1. System and Description

The experimental plant is a lab-scale quadruple tank system composed by Quanser coupled tanks [18] is a laboratory scale. This multivariable hydraulic process (see Figure **??** is composed of:

- Four interconnected water tanks, each of them is fed by some input flow rate. Each of the four tanks has a maximum water level in the range of 30 cm and is equipped with discharge valve and a differential pressure sensor to measure the gauge pressure at the tank base, which is proportional to the water column height above it via (3.1), and outputs an analog voltage converted to a level measurement through the affine calibration relation (3.6);
- Two pumps that are driven by voltage commands and deliver flow rates to tanks;
- Two three-way valves that split the incoming flow and allow to configure the coupling between the manipulated inputs and the tank levels. For pump 1, a fraction γ_1 of the flow is directed to tank 3 while the remaining fraction $1 - \gamma_1$ is directed to tank 2. For pump 2, a fraction γ_2 is directed to tank 1 while $1 - \gamma_2$ is directed to tank 4. Tanks 1 and 3 drain into tanks 2 and 4 respectively, creating additional internal coupling and a multi stage flow path (see Fig. 3.1).

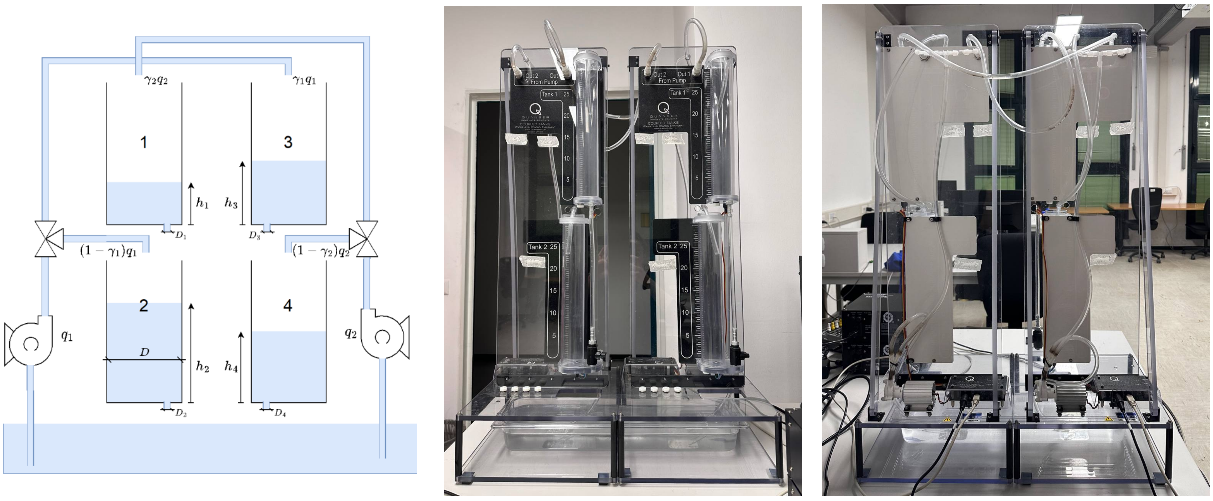

(a) QTS schematic. (b) QTS front view. (c) QTS back view.

Figure 3.1: QTS schematic and lab-scale plant.

The notation and physical parameters used to model the system are collected in Table 3.1. Geometric parameters such as D and S are fixed by the physical construction of the plant, while the valve outlet areas A_i and the pump and split maps are obtained from the grey-box identification procedure described in Section 3.4. The coefficients A_i are the primary source of parametric uncertainty in the model, since they are subject to drift with operating conditions and component ageing; managing this uncertainty is one of the central themes of this thesis.

Table 3.1: Quadruple-tank notation and identified parameters.

Notation	Description	Units
$h_i,\ i \in \{1,\dots,4\}$	Level in the i-th tank	cm
$v_i,\ i \in \{1,\dots,4\}$	Sensor output voltages (measured)	V
$\alpha_i,\ i \in \{1,\dots,4\}$	Sensor gain (calibration)	cm/V
$\beta_i,\ i \in \{1,\dots,4\}$	Sensor bias (calibration)	cm
$q_j,\ j \in \{1,2\}$	Pump flow rates	cm^3/s
$V_j,\ j \in \{1,2\}$	Pump command voltages (input)	V
$\gamma_j \in (0,1),\ j \in \{1,2\}$	Pump split coefficients	—
D	Diameter of the tank	cm
$D_i,\ i \in \{1,\dots,4\}$	Diameter of the discharge valves	cm
$S = 0.25\,\pi D^2$	Tank cross-sectional area	cm^2
$A_i = 0.25\,\pi D_i^2$	Valve cross-sectional area	cm^2
g	Gravitational acceleration	cm/s^2

The controlled outputs are the water levels in the bottom tanks,

$$y(t) = \begin{bmatrix} h_2(t) \\ h_4(t) \end{bmatrix}, \qquad 0.1 \le h_i \le 27.0 \text{ cm},$$

while the manipulated inputs are the pump voltages,

$$u(t) = \begin{bmatrix} V_1(t) \\ V_2(t) \end{bmatrix}, \qquad 4 \le V_j \le 18 \text{ V},$$

subject to a rate constraint on the voltage increments,

$$|\Delta V_j| \le 0.5 \text{ V}.$$

The level bounds reflect the physical tank geometry, while the voltage and rate limits are imposed to protect the pump hardware and avoid aggressive actuation.

3.2. Sensing, Actuation and Data Acquisition Hardware

Each tank is equipped with an analog level sensor which provides an estimation of the level based on the deformation of a membrane due to the pressure differential

$$\Delta P = \rho g h, \tag{3.1}$$

where ρ $[kg/cm^3]$ is the water density, $g = 981$ cm/s^2 is the gravitational acceleration, and h $[cm]$ is the water column height above the membrane. In practice, the sensor output heavily depends on ambient temperature and humidity (slow drift) and water temperature (drift during operations) [40]. For that reason, the calibration procedure used to correct for these effects is described in Section 3.5.

Sensor signals are routed to the data acquisition hardware using shielded cables to minimize electrical interference. Physical graduated scales (ruler markings) are mounted on the front panel alongside each tank, allowing the operator to read the water level directly by eye and cross-check it against the sensor voltage during experiments to detect gross calibration errors or sensor faults.

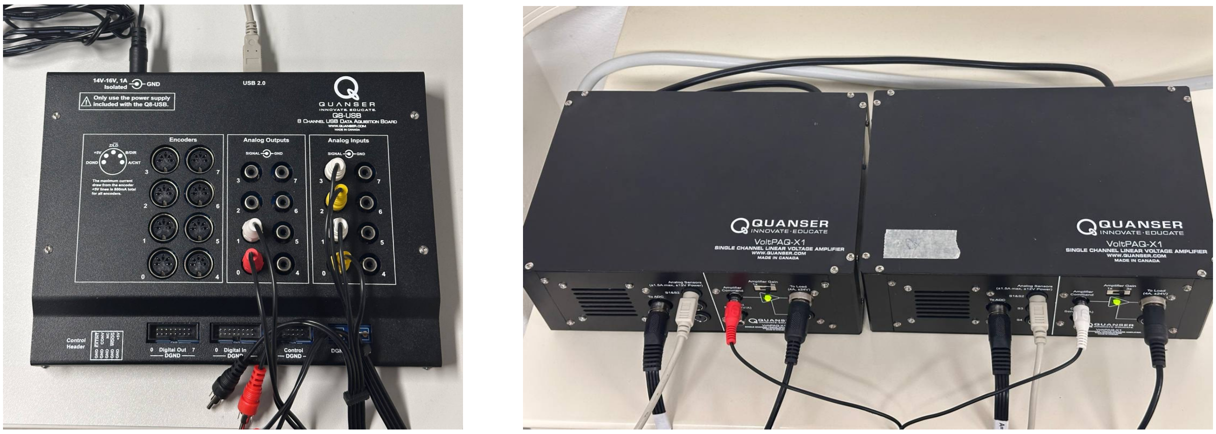

(a)ADC/DAC. (b) Power supplies.

Figure 3.2: Hardware details.

ADC/DAC (see Fig. 5.26) converts

- Sensors' analog signals to digital via USB
- Actuators' digital commands to analog signals for the power supply

Power supplies

- Apply the voltage (0–24V) to the pumps
- Condition the differential pressure sensor voltage
- Ensure the equipment's insulation from the grid

The actuators are modeled in a grey-box form that captures the main static nonlinearities of the laboratory hardware [40]. Each pump is assumed to behave as a volumetric device whose delivered flow rate depends on the applied command voltage V_j. A voltage-dependent gain $\kappa_j(\cdot)$ is introduced to account for dead-zone and nonidealities at low actuation levels,

$$q_j = \kappa_j(V_j)\, V_j, \tag{3.2}$$

where $\kappa_j(\cdot)$ is identified from experimental data. The pump flow is then routed through a three-way valve that splits q_j into two branches. In an ideal valve the split ratio would be fixed by the valve geometry; however, at low pump voltages the reduced flow velocity causes the actual split to deviate from its nominal value due to friction and gravity effects within the valve body. A voltage-dependent coefficient $\gamma_j(V_j) \in (0, 1)$ is therefore introduced to capture this behaviour,

$$q_j^{(1)} = \gamma_j(V_j)\, q_j, \qquad q_j^{(2)} = \big(1 - \gamma_j(V_j)\big)\, q_j, \tag{3.3}$$

where $q_j^{(1)}$ and $q_j^{(2)}$ denote the flow directed to the upper and lower tanks respectively, and $\gamma_j(\cdot)$ is identified from experimental data alongside $\kappa_j(\cdot)$. This parametrization allows the model to reproduce the observed loss of actuation effectiveness near the lower operating range and improves prediction accuracy for constrained control design.

Correct wiring of all cables is essential for laboratory testing, since even minor connection errors can corrupt measurements, prevent actuation, or introduce intermittent faults that invalidate an experiment. For this reason, before each run the cabling should be checked for consistency against the wiring diagram, including sensor polarity, grounding, channel assignments, and actuator connections, so that changes in closed-loop behaviour can be attributed to the controller rather than to the setup.

3.3. First-principle Modeling

Outflow through the discharge valve of tank $i = 1, ..., 4$ follows Torricelli's law,

$$q_{\mathrm{out},i}(t) = A_i \sqrt{2g\, h_i(t)}, \tag{3.4}$$

Combining the Torricelli outflows (3.4), the inter-tank gravity drainage and the pump inflows (3.2)-(3.3), the four mass balances of the system shown in Figure 3.1 read as follows:

$$\begin{aligned} S\,\dot{h}_1 &= -A_1\sqrt{2g\,h_1} + \kappa_2(V_2)\,\gamma_2(V_2)\,V_2, && \text{(3.5a)}\\ S\,\dot{h}_2 &= \quad A_1\sqrt{2g\,h_1} - A_2\sqrt{2g\,h_2} + \kappa_1(V_1)\big(1-\gamma_1(V_1)\big)V_1, && \text{(3.5b)}\\ S\,\dot{h}_3 &= -A_3\sqrt{2g\,h_3} + \kappa_1(V_1)\,\gamma_1(V_1)\,V_1, && \text{(3.5c)}\\ S\,\dot{h}_4 &= \quad A_3\sqrt{2g\,h_3} - A_4\sqrt{2g\,h_4} + \kappa_2(V_2)\big(1-\gamma_2(V_2)\big)V_2. && \text{(3.5d)} \end{aligned}$$

Note that, cross-coupling is the source of the multivariable interaction that makes the control problem nontrivial.

The water levels are related to the level sensor voltages v_i by affine calibration relations,

$$h_i = \alpha_i\, v_i + \beta_i, \qquad i = 1, \ldots, 4, \tag{3.6}$$

where α_i and β_i are the sensor gain and bias identified through the calibration procedure described in Section 3.5.

Defining the state vector $x(t) = \begin{bmatrix} h_1(t) & h_2(t) & h_3(t) & h_4(t) \end{bmatrix}^\top$ and the input vector $u(t) = \begin{bmatrix} V_1(t) & V_2(t) \end{bmatrix}^\top$, the system can be written in the standard nonlinear state-space form

$$\dot{x}(t) = f\big(x(t),\, u(t)\big), \qquad y(t) = C\,x(t), \tag{3.7}$$

where the output matrix C selects the controlled levels in tanks 2 and 4,

$$C = \begin{bmatrix} 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix}. \tag{3.8}$$

For controller design it is convenient to compute an equilibrium $(\bar{x},\, \bar{u})$ corresponding to a desired output reference $\bar{y}$. The equilibrium satisfies $f(\bar{x},\, \bar{u}) = 0$ and $\bar{y} = C\bar{x}$. Because $\kappa_j(\cdot)$ and $\gamma_j(\cdot)$ are stored as numerical maps, the equilibrium is found by solving a constrained nonlinear problem: stacking the decision vector $z = [\,\bar{x}^\top \;\; \bar{u}^\top\,]^\top$ and defining

$$g(z) = \begin{bmatrix} f(\bar{x},\, \bar{u}) \\ \bar{y} - C\bar{x} \end{bmatrix}, \tag{3.9}$$

the problem $g(z) = 0$ is solved subject to box constraints $z_{\min} \leq z \leq z_{\max}$ using a constrained nonlinear solver.

Using perturbation variables $\delta x = x - \bar{x}$ and $\delta u = u - \bar{u}$, the continuous-time linearized model is

$$\delta\dot{x}(t) = A_c\,\delta x(t) + B_c\,\delta u(t), \qquad \delta y(t) = C_c\,\delta x(t), \tag{3.10}$$

where

$$A_c = \left.\frac{\partial f}{\partial x}\right|_{(\bar{x},\,\bar{u})}, \qquad B_c = \left.\frac{\partial f}{\partial u}\right|_{(\bar{x},\,\bar{u})}. \tag{3.11}$$

Since the pump and valve maps are stored as numerical lookup tables, the Jacobians are computed using central finite differences.

The discrete-time model with sampling time T_s is obtained via Forward Euler discretization,

$$\delta x_{k+1} = A_d\,\delta x_k + B_d\,\delta u_k, \qquad \delta y_k = C_c\,\delta x_k, \tag{3.12}$$

where $A_d = I + T_s A_c$ and $B_d = T_s B_c$. This approximation is adequate for the sampling time $T_s = 1\,\mathrm{s}$ used in this work.

The choice is motivated by the standard rule of thumb for discrete-time process control is to select T_s such that the open-loop settling time contains between 5 and 50 samples [1]. The dominant open-loop time constants of the quadruple tank, identified from the linearized model at the operating point, are on the order of 30–100 s, yielding an open-loop 99% settling time of approximately 50–60 s. Choosing $T_s = 1\,\mathrm{s}$ places roughly 50 samples within this window, satisfying the upper end of the guideline while respecting the hardware constraint that all real-time computation must complete within one sample. The choice is further supported by the experimental observation that the closed-loop settling times under MPC are 25–50 s (see Chapter 5), giving 25–50 samples per transient, well within the recommended range.

3.4. Grey-box Nonlinear Modeling

To reduce the gap between controller design and experimental validation, the thesis relies on a provided grey-box nonlinear model that combines first principles structure with identified parameters and numerical maps. The model includes: mass balances for each tank, Torricelli outflow relations, pump voltage to flow characteristics, valve split coefficients, actuator delays and sensor calibration.

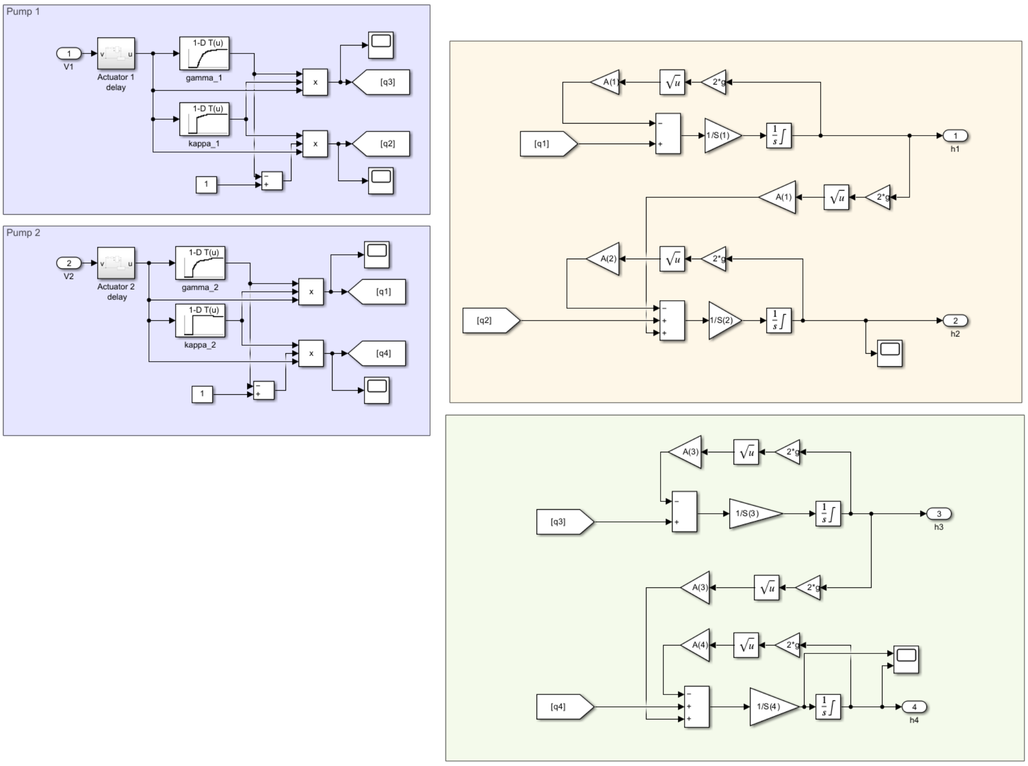

Figure 3.3: Grey-box model including nonlinear hydraulic balances, numerical pump and valve maps, actuator delays and sensing blocks.

This grey-box model is used in two ways. First, it serves as a high fidelity quadruple tank simulator environment for nonlinear closed-loop tests. Second, it is used to compute equilibria and linearizations required by LQI and by the linear MPC designs. The identified parameters are listed in Table 3.1.

3.5. Sensor Calibration

Accurate level measurements are essential for both model identification and controller performance. The quadruple tank system is equipped with differential pressure sensors that provide analog voltage signals proportional to the water level in each tank. In practice, these measurements are affected by sensor gain mismatch, bias offsets and slow drifts caused by ambient conditions and water temperature. For this reason, a dedicated calibration procedure was carried out prior to controller design.

Following the standard approach described in the Politecnico di Milano laboratory notes [40], each level sensor is modeled using an affine relation (3.6).

The calibration procedure was performed independently for each tank following these steps:

1. Close the outlet valve of the tank to isolate it hydraulically.
2. Select a set of known reference levels within the operating range of the tank.
3. For each reference level, fill the tank manually until the desired height is reached, as verified using the graduated scale on the tank column.
4. Wait until steady conditions are achieved and record the corresponding sensor voltage v_i.
5. Estimate the gain α_i and bias β_i by solving a linear least squares problem that minimises the squared error between the measured reference levels and the affine sensor model (3.6).

This procedure was repeated for all four tanks. The identified values of α_i and β_i were subsequently updated in the control software and used consistently in all experiments presented in this thesis.

Although this calibration method is affected by small inaccuracies due to parallax error and residual sensor drift, it significantly improves the consistency between measured and actual water levels and reduces steady-state offsets observed during closed-loop operation.

3.6. Implementation

All controllers are implemented in MATLAB and Simulink and executed on the laboratory workstation using Simulink Desktop real-time. Sensor voltages are acquired through the ADC interface, then converted into level measurements through calibration and filtering. Pump voltages are applied through the DAC interface. The main control loop runs with sampling time $T_s = 1$ s, while signal conditioning and logging run at a higher internal rate.

Safety is enforced through software and hardware limits. In particular, voltage saturation bounds are imposed on pump commands and level limits are monitored to prevent overflow. An emergency stop procedure allows manual intervention.

All controllers are tested using the following experimental procedure to ensure a fair comparison:

1. Initialization: bring the system close to the chosen equilibrium and wait until levels settle.
2. Reference tracking: apply step changes to the two water level references and record the transient response.
3. Disturbance rejection: introduce repeatable disturbances, then evaluate recovery.
4. Data logging: record time, measured levels, references and applied voltages for offline analysis.
5. Metric computation: compute tracking and effort metrics.

Between tests, tanks are emptied to have the same start-up conditions. The collected datasets form the basis for the comparative analysis and for discussing practical trade offs between LQI and MPC variants.

4 | Optimal Control Strategies

This chapter presents the design and tuning of the five control strategies applied to the quadruple tank system. The controllers are organised in order of increasing complexity, each addressing limitations of its predecessor. The Linear Quadratic Regulator with integral action (LQI) serves as a baseline, providing unconstrained optimal feedback with offset-free tracking. Nominal MPC extends this by enforcing explicit constraints on tank levels and pump voltages within a finite-horizon optimization. MPC with successive linearisation (SL-MPC) further improves prediction accuracy by updating the linear model at each sampling instant to account for the nonlinear Torricelli outflow dynamics. Scenario-based MPC (SC-MPC) introduces robustness to parametric uncertainty by enforcing constraints across multiple sampled realisations of the valve coefficients. Finally, adaptive SC-MPC closes the loop on uncertainty by combining online set-membership identification with zonotopic parameter bounding, progressively reducing conservatism as experimental data accumulates.

A common experimental procedure is shared across all controllers: the same sampling time, actuator limits, signal conditioning and reference trajectories are retained so that performance differences can be attributed to the control architecture rather than to implementation choices. The key performance indicators used for comparison are introduced in Chapter 5.

4.1. Linear Quadratic Regulator with Integral Action

This section explains the LQI method for the quadruple tank system which serves as the baseline against which the MPC architectures.

The controller uses the full state vector $x(t) = [h_1(t) \quad h_2(t) \quad h_3(t) \quad h_4(t)]^\top$ with manipulated inputs $u(t) = [V_1(t) \quad V_2(t)]^\top$. The regulated outputs are the levels in tanks 2 and 4,

$$y(t) = \begin{bmatrix} h_2(t) \\ h_4(t) \end{bmatrix} = C\,x(t), \qquad C = \begin{bmatrix} 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix}. \tag{4.1}$$

For a constant reference $\bar{y} = [\bar{r}_2 \;\; \bar{r}_4]^\top$, the objective is to regulate $y(t) \to \bar{y}$ with no steady-state offset under model-plant mismatch and constant disturbances. Pure LQR does not guarantee offset-free tracking; but the augmentation with integral action addresses this limitation.

The system model is linearized around the equilibrium $(\bar{x}, \bar{u})$ associated with the target output $\bar{y}$. The equilibrium satisfies

$$0 = f(\bar{x}, \bar{u}), \qquad \bar{y} = C\bar{x},$$

where $f(\cdot)$ denotes the nonlinear grey-box dynamics introduced in Chapter 3. Because the pump flow functions $\kappa_j(\cdot)$ and the split ratio functions $\gamma_j(\cdot)$ are implemented as numerical lookup tables rather than closed-form expressions, the equilibrium is computed numerically as a constrained nonlinear feasibility problem.

The linearized continuous-time model (3.10) with Jacobians (3.11) evaluated at the operating point $(\bar{x}, \bar{u})$ is discretized with sampling time $T_s = 1$ s via zero-order hold (ZOH), yielding the discrete-time model (3.12). Note that the first definition used a Forward Euler approximation; here ZOH discretization is adopted instead to improve accuracy of the state-space matrices used in the LQI design.

Following the LQI augmentation from Section 2.2, the plant state is extended with the integrated output error. In deviation coordinates, the output error is $\delta e_k = \delta r_k - \delta y_k$ where $\delta r_k = r_k - \bar{y}$ and $\delta y_k = C_d \, \delta x_k + D_d \, \delta u_k$. The integral state $x_{I,k} \in \mathbb{R}^2$ evolves as

$$x_{I,k+1} = x_{I,k} + T_s \delta e_k.$$

Stacking the deviation state and the integral state into

$$\tilde{x}_k = [\delta x_k^\top \;\; x_{I,k}^\top]^\top \in \mathbb{R}^6, \tag{4.2}$$

the augmented discrete-time dynamics are

$$\tilde{x}_{k+1} = \underbrace{\begin{bmatrix} A_d & 0 \\ -C_d \, T_s & I \end{bmatrix}}_{\tilde{A}} \tilde{x}_k + \underbrace{\begin{bmatrix} B_d \\ -D_d \, T_s \end{bmatrix}}_{\tilde{B}} \delta u_k + \underbrace{\begin{bmatrix} 0 \\ T_s \, I \end{bmatrix}}_{\tilde{E}} \delta r_k \tag{4.3}$$

Only the pair $(\tilde{A}, \tilde{B})$ is required for the gain computation; the reference term is handled separately at implementation level through the integral state update.

The LQI gain is obtained by solving the Discrete Algebraic Riccati Equation (DARE) for the augmented pair $(\tilde{A},\, \tilde{B})$ with cost matrices $\tilde{Q}$ and $\tilde{R}$ chosen according to Bryson's rule as discussed in Section 2.2. The specific numerical choices for the quadruple tank system are given below.

For each tank, the maximum acceptable level deviation from the operating point is set to $x_{\max} = 0.7\,\text{cm}$, yielding

$$Q_x = \operatorname{diag}\left(\frac{1}{0.7^2},\ \frac{1}{0.7^2},\ \frac{1}{0.7^2},\ \frac{1}{0.7^2}\right). \tag{4.4}$$

The value $x_{\max} = 0.7\,\text{cm}$ was selected based on two considerations. First, it represents approximately 5–8% of the feasible operating range of 9–15 cm, which is a reasonable tolerance for level regulation in a water process where sensor calibration errors are on the order of 0.1–0.2 cm and measurement noise fluctuations are present.Second, it is large enough to avoid exciting measurement noise aggressively yet small enough to produce a clearly visible correction signal: halving $x_{\max}$ to 0.35 cm was found experimentally to cause pump saturation immediately after setpoint changes, while doubling it to 1.4 cm produced sluggish responses with visible steady-state offset before the integrator compensated. The integral states are weighted more softly with $x_{I,\max} = 2\,\text{cm·s}$,

$$Q_I = \operatorname{diag}\left(\frac{1}{2^2},\ \frac{1}{2^2}\right), \tag{4.5}$$

so that the full state weighting matrix is

$$\tilde{Q} = \begin{bmatrix} Q_x & 0 \\ 0 & Q_I \end{bmatrix}. \tag{4.6}$$

The value $x_{I,\max} = 2\,\text{cm·s}$ reflects the typical magnitude of the accumulated tracking error during a settling transient: with experimental settling times of 25–50 s and a steady-state error that decays from roughly 1 cm to zero, the integrated error is on the order of 1–2 cm·s. Weighting the integrator state with this bound ensures it contributes meaningfully to the cost without dominating the state penalty. The ratio $x_{I,\max}/x_{\max} \approx 2.9$ ensures that the integrator builds up gradually, preventing aggressive transients while still eliminating steady-state offset within a reasonable time [1].

The pump voltage span is $u_{\text{span}} = u_{\max} - u_{\min} = 18 - 4 = 14\,\text{V}$. The input weight is

$$\tilde{R} = \alpha \operatorname{diag}\left(\frac{1}{u_{\text{span}}^2},\ \frac{1}{u_{\text{span}}^2}\right), \qquad \alpha = 100. \tag{4.7}$$

The large scaling factor $\alpha = 100$ penalizes aggressive voltage variations. This choice produces smooth input signals, which is desirable for pump longevity and for establishing a fair comparison baseline: the MPC controllers are tuned with consistent weight philosophy so that performance differences can be attributed to the control architecture.

The LQI gain $\tilde{K} = \begin{bmatrix} K_x & K_I \end{bmatrix}$ is computed by solving the DARE with $(\tilde{A}, \tilde{B}, \tilde{Q}, \tilde{R})$ and the control law in deviation coordinates is

$$\delta u_k = -K_x\, \delta x_k - K_I\, x_{I,k}, \qquad u_k = \bar{u} + \delta u_k. \tag{4.8}$$

Pump voltages are constrained to the hardware safe range

$$u_k \in [u_{\min}, u_{\max}]^2, \qquad u_{\min} = 4\text{ V}, \quad u_{\max} = 18\text{ V}. \tag{4.9}$$

The applied input is $u_k^{\text{sat}} = \text{sat}(u_k)$ and state propagation uses the saturated deviation $\delta u_k^{\text{sat}} = u_k^{\text{sat}} - \bar{u}$. Unlike MPC, this projection-based saturation does not anticipate future constraint activity [2].

To prevent integrator windup during prolonged saturation, the integral state is clamped after each update,

$$x_{I,k} \in \left[-I_{\text{clamp}},\ I_{\text{clamp,upper}}\right]^2, \qquad I_{\text{clamp}} = 2, \qquad I_{\text{clamp,upper}} = u_{\max}. \tag{4.10}$$

This limits the maximum stored integral energy and prevents prolonged overshoot once the actuator exits saturation. More advanced anti-windup strategies such as back-calculation were considered but not needed given the moderate transients in the experimental protocol.

On the real plant a two-stage initialization ensures a repeatable and safe start. In the first stage, the equilibrium voltage $\bar{u}$ is applied as a constant open-loop input, allowing the tank levels to approach the intended operating region without triggering integrator windup or actuator saturation. In the second stage, the controller is engaged but integral action is disabled for the first 15 s, preventing the integral state from accumulating a large value during the settling transient. After this window, the LQI loop is fully enabled. This procedure reduces variability across experimental runs and is applied identically when switching between controllers.

The LQI controller is implemented in a discrete-time Simulink subsystem block. The measured levels are converted to deviation coordinates using the stored equilibrium $\bar{x}$.

The output error is computed as the difference between the reference and the measured regulated outputs (h_2, h_4), integrated with gain T_s and optional gating for the initialization window. The control action follows (4.8), with the absolute command $u_k = \bar{u} + \delta u_k$ saturated to (4.9) before being sent to the plant.

Table 4.1: LQI design parameters.

Parameter	**Symbol**	**Value**
Operating point output	$\bar{y}$	$[10,\ 10]^\top$ cm
Sampling time	T_s	1 s
Maximum level deviation	$x_{\max}$	0.7 cm
Maximum integral deviation	$x_{I,\max}$	2 cm·s
Input weight scaling	α	100
Voltage bounds	$[u_{\min},\ u_{\max}]$	[4, 18] V
Initialization delay	—	15 s

The LQI controller provides a computationally efficient baseline with systematic tuning and offset-free tracking. Its primary advantage is the absence of online optimization: the gain is computed once offline from the DARE solution. Its primary limitations are the lack of explicit constraint handling, which is enforced only through post-hoc saturation, and its reliance on a single linearization point. These limitations motivate the transition to MPC, where constraints enter the optimization formulation directly [2]. The values of the weight matrices Q and R are reused as a reference point for MPC tuning.

4.2. Model Predictive Control

This section develops two MPC controllers for the quadruple tank process. The first is a nominal linear MPC designed around a fixed operating point. The second is an MPC with successive linearization, where the prediction model is updated online at each sampling instant to compensate for the level-dependent gain variations introduced by the Torricelli outflow nonlinearity. Both controllers are implemented in discrete time with $T_s = 1$ s, include explicit constraint handling and follow the same experimental protocol as the LQI baseline to ensure a fair comparison.

4.2.1. Nominal velocity-form MPC

To achieve offset-free tracking, the MPC is formulated with explicit internal action and in velocity form as described in Section 2.3. The incremental quantities are $\Delta x_k = x_k - x_{k-1}$

and $\Delta u_k = u_k - u_{k-1}$, and the tracking error is

$$e_k = y_k - \bar{y}_k, \tag{4.11}$$

where the sign convention of $y_k - \bar{y}_k$ matches the Simulink implementation. The augmented velocity-form state is

$$\xi_k = \begin{bmatrix} \Delta x_k \\ e_k \end{bmatrix} \in \mathbb{R}^{\tilde{n}}, \quad \tilde{n} = n_x + n_y = 6, \tag{4.12}$$

and the dynamics of the enlarged system is

$$\xi_{k+1} = \tilde{A}\,\xi_k + \tilde{B}\,\Delta u_k, \qquad \tilde{A} = \begin{bmatrix} A_d & 0 \\ C_d A_d & I \end{bmatrix}, \quad \tilde{B} = \begin{bmatrix} B_d \\ C_d B_d \end{bmatrix}. \tag{4.13}$$

In the Simulink implementation, Δx_k is computed from the current measured state x_k and a stored previous value x_{k-1} using a persistent memory variable updated each sample.

Following the guideline in [1] that N should be chosen to cover the settling time of the process, the prediction horizon is set to $N = 60$ steps, corresponding to 60 s at $T_s = 1$ s. This approximately matches the open-loop settling time of the slower tank channels, ensuring that the optimization looks far enough ahead to anticipate constraint activation and that the terminal cost P provides an adequate approximation of the infinite-horizon tail [2]. Increasing N beyond 60 steps yields negligible improvement in closed-loop performance because the predicted trajectories have essentially settled to steady state, while the number of decision variables grows as $n_u N$, increasing QP solve time [1].

The control horizon is set to $N_c = 10$ steps, so that only the first N_c input increments are optimised freely and the remaining $N - N_c = 50$ are held constant ($\Delta u = 0$ for $i \geq N_c$). This blocking strategy reduces the number of decision variables from $n_u N = 120$ to $n_u N_c = 20$, which is essential for real-time feasibility at $T_s = 1$ s [1]. The ratio $N/N_c = 6$ retains sufficient degrees of freedom to handle the cross-coupling between the four tanks and to reject step disturbances, while the constant-input tail over the remaining 50 steps does not degrade performance since the dominant transient is captured within the first N_c steps.

Stacking the predicted states and decision variables as

$$\Xi = \begin{bmatrix} \xi_{k+1} \\ \vdots \\ \xi_{k+N} \end{bmatrix} \in \mathbb{R}^{\tilde{n}N}, \qquad \Delta U = \begin{bmatrix} \Delta u_k \\ \vdots \\ \Delta u_{k+N_c-1} \end{bmatrix} \in \mathbb{R}^{n_u N_c},$$

the prediction model is written in matrix form $\Xi = S_x\,\xi_k + S_u\,\Delta U$. For nominal MPC, the matrices S_x and S_u are precomputed offline.

The cost function penalizes augmented state deviation and control increments,

$$J = \Xi^\top \bar{Q}\,\Xi + \Delta U^\top \bar{R}\,\Delta U = \sum_{i=1}^{N-1} \xi_{k+i}^\top Q\,\xi_{k+i} + \xi_{k+N}^\top P\,\xi_{k+N} + \sum_{i=0}^{N_c-1} \Delta u_{k+i}^\top R\,\Delta u_{k+i}, \quad (4.14)$$

where $Q = \begin{bmatrix} Q_x & 0 \\ 0 & Q_e \end{bmatrix}$ uses the same Bryson scaling as the LQI design and the terminal weight P is computed by solving $(\tilde{A}, \tilde{B}, Q, R)$, approximating the infinite-horizon tail cost as discussed in Section 2.3.

Substituting the prediction model into (4.14) yields the QP problem

$$\min_{\Delta U} \ \tfrac{1}{2}\,\Delta U^\top H\,\Delta U + f^\top \Delta U, \quad (4.15)$$

with $H = 2(S_u^\top \bar{Q}\,S_u + \bar{R})$ and $f = 2S_u^\top \bar{Q}\,S_x\,\xi_k$. The Hessian matrix H is positive definite and precomputed offline; the gradient vector f is evaluated online as $f = M_\xi\,\xi_k$ where $M_\xi = 2S_u^\top \bar{Q}\,S_x$.

Three classes of inequality constraints are enforced:

- Slew rate constraints: each control increment is bounded elementwise,

$$|\Delta u_{k+i}| \le \Delta u_{\max} = 0.5\ \text{V}, \qquad i = 0, \ldots, N_c - 1. \quad (4.16)$$

- Absolute input saturation: the absolute input must satisfy

$$u_{\min} \le u_{k+i} \le u_{\max}, \qquad u_{\min} = 4\ \text{V}, \quad u_{\max} = 18\ \text{V}. \quad (4.17)$$

  The absolute input sequence is reconstructed from increments as $u_{k+i} = u_{k-1} + \sum_{j=0}^{i} \Delta u_{k+j}$. Defining L as the lower-triangular summation matrix that maps ΔU to stacked absolute inputs, this yields linear inequalities whose right-hand side depends on u_{k-1} and is recomputed online.

- Tank level constraints:

$$h_{\min} \leq h_i(k+j) \leq h_{\max}, \qquad h_{\min} = 0.1 \text{ cm}, \quad h_{\max} = 27.0 \text{ cm}, \tag{4.18}$$

  for all four tanks over the prediction horizon. Predicted levels are extracted from the augmented state using the extraction matrix $E = [I_{n_x} \ \ 0]$ and propagated through the prediction matrices, yielding linear constraints on ΔU with time-varying right-hand sides depending on the current measured state x_k.

At each sampling instant the MPC algorithm performs:

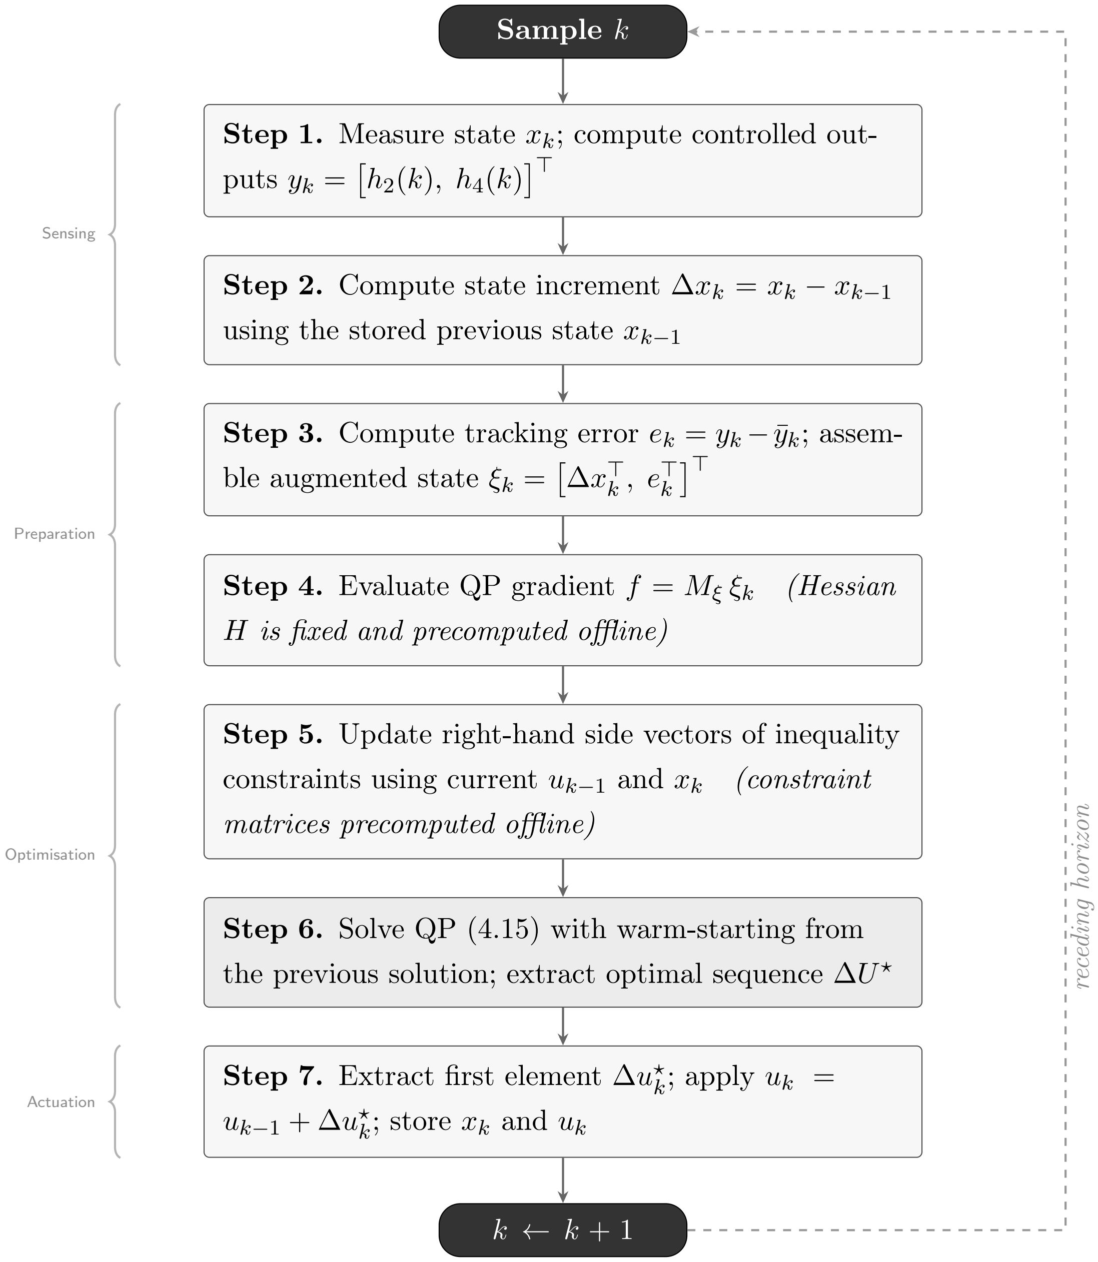

Figure 4.1: Nominal MPC execution loop at each sampling instant k. Left annotations indicate the logical phase of each step.

If the QP solver fails to converge, a fallback strategy applies $\Delta u_k = 0$ to hold the previous control action.

4.2.2. MPC with successive linearization

While nominal MPC provides good performance near the design operating point, prediction accuracy degrades over wide operating ranges because the fixed Linear Time-

Invariant (LTI) model does not capture the level-dependent hydraulic dynamics. For that reason, successive linearization can improve accuracy by recomputing the Jacobians at each sample using the current measured state. This approach is sometimes referred to as quasi-LPV or frozen-Jacobian MPC in the literature, since the linearization is updated between samples but held constant over the prediction horizon.

The nonlinear model is linearized around the current state x_k and the previous control input u_{k-1}. From the mass balances (3.5), linearizing the Torricelli outflow $q_{\text{out},i} = A_i\sqrt{2g\,h_i}$ around the current level yields:

$$\frac{\partial}{\partial h_i}\left(A_i\sqrt{2g\,h_i}\right) = \frac{A_i\sqrt{g}}{\sqrt{2\,h_i}} \triangleq a_i(h_i). \tag{4.19}$$

where valve's cross sections A_i are the identified parameters. To avoid singularities, measured levels are clamped to $h_i \leftarrow \max(h_i,\, 0.1)$ before evaluation.

The resulting state-dependent continuous-time matrix has the structure:

$$A_c(x_k) = \frac{1}{S}\begin{bmatrix} -a_1 & 0 & 0 & 0 \\ a_1 & -a_2 & 0 & 0 \\ 0 & 0 & -a_3 & 0 \\ 0 & 0 & a_3 & -a_4 \end{bmatrix}, \tag{4.20}$$

Additionally, the input matrix $B_c(u_{k-1})$ depends on the pump flow gains and split ratios evaluated at the previously applied voltage,

$$B_c(u_{k-1}) = \frac{1}{S}\begin{bmatrix} 0 & \kappa_2\,\gamma_2 \\ \kappa_1(1-\gamma_1) & 0 \\ \kappa_1\,\gamma_1 & 0 \\ 0 & \kappa_2(1-\gamma_2) \end{bmatrix}, \tag{4.21}$$

where maps $\kappa_j(\cdot)$, $\gamma_j(\cdot)$ are evaluated at $V_j(k-1)$ by linear interpolation on the identified lookup tables. Pump voltages are clamped into the interval $[4,\, 18]\,\mathrm{V}$ before interpolation to remain within the valid range.

At each sample, the continuous-time linearized matrices (4.20)–(4.21) are discretized using Forward Euler,

$$A_d(k) = I + A_c(x_k)\,T_s, \qquad B_d(k) = B_c(u_{k-1})\,T_s. \tag{4.22}$$

The augmented model in velocity form is then constructed online by matrices:

$$A_{\text{aug}}(k) = \begin{bmatrix} A_d(k) & 0 \\ C_d\, A_d(k) & I \end{bmatrix}, \quad B_{\text{aug}}(k) = \begin{bmatrix} B_d(k) \\ C_d\, B_d(k) \end{bmatrix}, \tag{4.23}$$

and the prediction matrices $S_x(k)$, $S_u(k)$ are recomputed. This implies that the Hessian matrix H, the gradient vector f and the constraint matrices of the QP problem are all updated at every sample,

$$H(k) = 2\big(S_u(k)^\top \bar{Q}\, S_u(k) + \bar{R}\big), \qquad f(k) = 2\, S_u(k)^\top \bar{Q}\, S_x(k)\, \xi_k. \tag{4.24}$$

The constraint structure is identical to nominal MPC, but the matrices involved in the level constraints now depend on $S_u(k)$ and $S_x(k)$ and are therefore updated online. The QP is solved using `quadprog` with warm starting as in the nominal case.

The terminal weight P is computed by solving the DARE only at the nominal operating point and reused across all operating conditions. While this does not provide a formal recursive feasibility guarantee under large model variations, it offers an adequate terminal cost approximation for the operating range considered in the experiments (9–15 cm). A state-dependent terminal cost recomputed at each sample would be theoretically more consistent but was not pursued due to the additional computational burden.

Both linear MPC controllers (the nominal and successive linearization variant) share identical tuning to maintain consistency with the LQI baseline. Table 4.2 summarises the main tuning parameters. The ratio $N/N_c = 6$ ensures sufficient look-ahead for the

Table 4.2: MPC tuning parameters.

Parameter	**Symbol**	**Value**
Sampling time	T_s	1 s
Prediction horizon	N	60
Control horizon	N_c	10
State weight (per tank)	$1/x_{\max}^2$	$1/0.7^2$
Error weight (per output)	$1/e_{\max}^2$	$1/2^2$
Input increment weight	R	$100 \cdot \mathrm{diag}(1/14^2,\ 1/14^2)$
Terminal weight	P	DARE on $(\tilde{A},\ \tilde{B},\ Q,\ R)$
Slew rate limit	$\Delta u_{\max}$	0.5 V
Voltage bounds	$[u_{\min},\ u_{\max}]$	[4, 18] V
Level bounds	$[h_{\min},\ h_{\max}]$	[0.1, 27.0] cm

slow tank dynamics (time constants on the order of 30–100 s) while keeping the number of decision variables at $n_u N_c = 20$, which is manageable for real-time execution every $T_s = 1\,\mathrm{s}$. After the control horizon, the input is held constant ($\Delta u = 0$ for $i \geq N_c$) along the prediction horizon.

The weight matrices use the same Bryson scaling as the LQI design: $x_{\max} = 0.7\,\mathrm{cm}$ for each tank level and $e_{\max} = 2\,\mathrm{cm}$ for each tracking error component, giving

$$Q_x = \operatorname{diag}\left(\tfrac{1}{x_{\max}^2}, \tfrac{1}{x_{\max}^2}, \tfrac{1}{x_{\max}^2}, \tfrac{1}{x_{\max}^2}\right), \qquad Q_e = \operatorname{diag}\left(\tfrac{1}{e_{\max}^2}, \tfrac{1}{e_{\max}^2}\right), \tag{4.25}$$

$$Q = \begin{bmatrix} Q_x & 0 \\ 0 & Q_e \end{bmatrix}. \tag{4.26}$$

The large value of R (see Table 4.2), combined with the slew rate constraint (4.16), produces smooth control signals consistent with the LQI baseline philosophy.

In simulations, both MPC controllers are tested on the nonlinear grey-box model using the same initial conditions and target references as the LQI baseline.

Successive linearization MPC, based on a model with $A_d(k)$ and $B_d(k)$, improves prediction accuracy because the matrices follow the current hydraulic regime through the level-dependent derivatives in (4.20) and the input-dependent pump maps in (4.21). The reference tracking remains consistent across a wider range of operating conditions, especially during large transients. The computational cost is higher than nominal MPC because the linearized model and constraints are reconstructed online at each sample.

Both controllers are also validated on the real plant using the lab setup described in Chapter 3. The system is first stabilized near the operating point using the open-loop equilibrium voltage, then reference changes are applied. Logged signals include tank levels, references, pump commands and internal controller variables. The experimental results and comparative analysis are presented in Chapter 5.

4.3. Scenario-Based Model Predictive Control

This section introduces a scenario-based MPC strategy to provide explicit guarantees on constraint satisfaction under parametric uncertainty. Building on the scenario approach framework described in Section 2.4, uncertainty model, scenario generation, constraint tightening and QP formulation are detailed for the quadruple tank plant.

A practical advantage of the SC-MPC is that it operates directly in closed-loop from the first sampling instant without requiring the open-loop stabilization phase used by LQI and the nominal MPC variants, since the robust tightening mechanism ensures safe constraint satisfaction from arbitrary initial conditions.

The dominant source of parametric uncertainty in the quadruple tank system is the outlet valve coefficients A_i, which directly affect the linearized outflow dynamics through Torricelli's law. Based on experimental identification, each coefficient is assumed to lie within

$$A_i \in [0.75\, A_i^{\text{id}},\;\; 0.90\, A_i^{\text{id}}], \tag{4.27}$$

corresponding to a relative uncertainty of -25% to -10% around the nominal identified values [40]. This asymmetric range reflects the practical observation that valve coefficients tend to decrease over time due to fouling and sediment accumulation.

The uncertainty set (4.27) is represented as a zonotope,

$$\mathcal{A} = A_{\text{center}} \oplus A_{\text{radius}} \cdot [-1,\, 1]^4, \tag{4.28}$$

with $A_{\text{center}} = 0.825\, A^{\text{id}}$ and $A_{\text{radius}} = 0.075\, A^{\text{id}}$.

Then, $K = 39$ scenarios are sampled uniformly from (4.28) and resampled independently at each sampling instant,

$$A_k^{(j)} = A_{\text{center}} + \xi_k^{(j)} \odot A_{\text{radius}}, \quad \xi_k^{(j)} \sim \mathcal{U}([-1, 1]^4), \quad j = 1, \ldots, K, \tag{4.29}$$

where the subscript k denotes that a fresh independent draw is taken at every time step. Since scenarios are resampled at every receding-horizon step rather than fixed offline, the closed-loop violation bound applies rather than the open-loop a priori bound, which justifies the use of the tighter formula $K = 2/\epsilon - 1$ [8]. The scenario count follows from the risk level computation (2.27) with $\rho = n_u = 2$ and $\epsilon = 5\%$, as derived in Section 2.4,

$$K = \frac{2}{0.05} - 1 = 39.$$

This ensures that output constraints are satisfied with at least 95% confidence under the assumed uncertainty model. Since scenarios are resampled at every receding-horizon step rather than fixed offline, the closed-loop violation bound applies rather than the open-loop a priori bound, which justifies the use of the tighter formula $K = 2/\epsilon - 1$ [8]. The choice of K balances the trade-off between robustness (more scenarios reduce violation probability) and computation (each additional scenario increases the constraint tightening cost) [8].

At each sampling instant, a nominal linear model is constructed using the same successive linearization procedure as before. The continuous-time Jacobians $A_c(x_k)$ and $B_c(u_{k-1})$ follow equations (4.20)–(4.21), discretized via Forward Euler. The velocity-form augmentation and prediction matrices S_x, S_u are recomputed online since the model changes at each sample. The cost function and QP structure are identical to the successive linearization MPC.

The central mechanism of SC-MPC is the enforcement of output constraints across all K scenarios. For each scenario $j = \{1, ..., K\}$, the sampled valve vector $A_i^{(j)}$ produces a scenario-specific system Jacobian,

$$A_c^{(j)}(x_k) = \frac{1}{S}\begin{bmatrix} -a_1^{(j)} & 0 & 0 & 0 \\ a_1^{(j)} & -a_2^{(j)} & 0 & 0 \\ 0 & 0 & -a_3^{(j)} & 0 \\ 0 & 0 & a_3^{(j)} & -a_4^{(j)} \end{bmatrix}, \quad a_i^{(j)} = \frac{A_i^{(j)}\sqrt{g}}{\sqrt{2\,h_i}}, \tag{4.30}$$

and the corresponding augmented matrix $A_{\text{aug}}^{(j)}$ is built using $A_d^{(j)} = I + A_c^{(j)} T_s$.

The free response (driven by system state ξ_k alone, without control action) is propagated for each scenario,

$$\xi_{k+i}^{(j),\text{free}} = \left(A_{\text{aug}}^{(j)}\right)^i \xi_k, \qquad i = 1, \ldots, N, j = \{1, ..., K\}. \tag{4.31}$$

An extraction matrix $E = [0_{n_y \times n_x} \;\; I_{n_y}]$ selects the error component from the augmented state. The output constraints are expressed in terms of the tracking error relative to the reference: $b_{\text{upper}} = h_{\max}\mathbf{1} - \bar{y}$ and $b_{\text{lower}} = h_{\min}\mathbf{1} - \bar{y}$.

The scenario-based tightening is applied exclusively to the controlled output constraints (tank levels h_2 and h_4), not to the input or slew-rate constraints (enforced directly on ΔU) and not to the uncontrolled tank levels h_1 and h_3. The latter are bounded only under the nominal prediction; their physical limits (0.1–27.0 cm) are sufficiently loose at the operating points considered that violations are not observed experimentally. Extending the scenario tightening to all four tank states would require augmenting the extraction matrix to $E = [I_{n_x} \;\; 0]$ at the cost of additional constraint rows in the QP, and is left as a possible refinement. Let E be the output selection matrix such that $y_{k+i} = E\,\xi_{k+i}$ extracts the controlled outputs from the augmented state. For each prediction step i, the worst-case output deviation across all K scenarios tightens the output bounds,

$$\begin{aligned} b_i^{\text{upper,tight}} &= \min_{j=1,\ldots,K}\left(b_{\text{upper}} - E\,\xi_{k+i}^{(j),\text{free}}\right), \\ b_i^{\text{lower,tight}} &= \max_{j=1,\ldots,K}\left(b_{\text{lower}} - E\,\xi_{k+i}^{(j),\text{free}}\right), \end{aligned} \tag{4.32}$$

where b_{upper} and b_{lower} collect the original output bounds over the horizon, and $\xi_{k+i}^{(j),\text{free}}$ is the free response of the augmented state under scenario j (zero nominal input). This tightening guarantees that the nominal prediction satisfying the tightened bounds also satisfies the original output bounds for all K scenarios [8].

The physical level bounds are further tightened by an Root Mean Square Error (RMSE) margin of 2.09 cm on each side to account for the gap between the linear prediction and the true nonlinear plant, yielding effective bounds $h_{\min} = 2.19\,\text{cm}$ and $h_{\max} = 24.91\,\text{cm}$ [40].

The SC-MPC is formulated as a QP problem. A slack variable $\sigma \geq 0$ is introduced to preserve feasibility when the tightened bounds are highly restrictive. In natural form the cost is

$$\min_{\Delta U,\sigma} \sum_{i=1}^{N-1} \xi_{k+i}^{\top} Q\, \xi_{k+i} + \xi_{k+N}^{\top} P\, \xi_{k+N} + \sum_{i=0}^{N_c-1} \Delta u_{k+i}^{\top} R\, \Delta u_{k+i} + \alpha\, \sigma^2, \tag{4.33}$$

where the last term is a quadratic penalty on the slack with $\alpha = 10^6$, ensuring that σ is driven to zero whenever the problem is feasible without relaxation. The quadratic form is consistent with the QP structure and avoids introducing a linear term in the objective. Substituting the prediction model $\Xi = S_x\, \xi_k + S_u\, \Delta U$ yields the equivalent compact QP form

$$\min_{\Delta U,\sigma} \frac{1}{2} \begin{bmatrix} \Delta U \\ \sigma \end{bmatrix}^{\top} \begin{bmatrix} H & 0 \\ 0 & 2\alpha \end{bmatrix} \begin{bmatrix} \Delta U \\ \sigma \end{bmatrix} + \begin{bmatrix} f \\ 0 \end{bmatrix}^{\top} \begin{bmatrix} \Delta U \\ \sigma \end{bmatrix}, \tag{4.34}$$

where H and f are computed from the online prediction matrices as in (4.24).

The slack variable enters the scenario-tightened output constraints as

$$S_{u,e}\, \Delta U - \sigma\, \mathbf{1} \leq b^{\text{upper,tight}}, \qquad -S_{u,e}\, \Delta U - \sigma\, \mathbf{1} \leq -b^{\text{lower,tight}}, \tag{4.35}$$

where $S_{u,e} = (I_N \otimes E)\, S_u$ extracts the error prediction. Slew rate and absolute input saturation constraints remain identical to (4.16)–(4.17). The complete constraint inequality is

$$\begin{bmatrix} A_{\text{slew}} & 0 \\ A_{\text{sat}} & 0 \\ S_{u,e} & -\mathbf{1} \\ -S_{u,e} & -\mathbf{1} \\ 0 & -1 \end{bmatrix} \begin{bmatrix} \Delta U \\ \sigma \end{bmatrix} \leq \begin{bmatrix} \Delta u_{\max}\, \mathbf{1} \\ u_{\max}\, \mathbf{1} - \mathbf{1} \otimes u_{k-1} \\ \mathbf{1} \otimes u_{k-1} - u_{\min}\, \mathbf{1} \\ b^{\text{upper,tight}} \\ -b^{\text{lower,tight}} \\ 0 \end{bmatrix}, \tag{4.36}$$

where the last row enforces $\sigma \geq 0$.

After solving the QP problem, the same K scenarios drawn at the current step are reused to propagate the predicted output trajectories under the optimal control sequence $\Delta U^{\star}$. The prediction envelope is

$$y_{\min}(i) = \min_{j=1,\ldots,K} y^{(j)}(i), \qquad y_{\max}(i) = \max_{j=1,\ldots,K} y^{(j)}(i),\ i = 1,\ldots,N$$

yielding a $4N = 240$ element vector $[\min h_2;\ \max h_2;\ \min h_4;\ \max h_4]$ logged for monitoring. These values provides post-experiment confirmation that the predicted uncertainty spread is consistent with the observed plant behaviour.

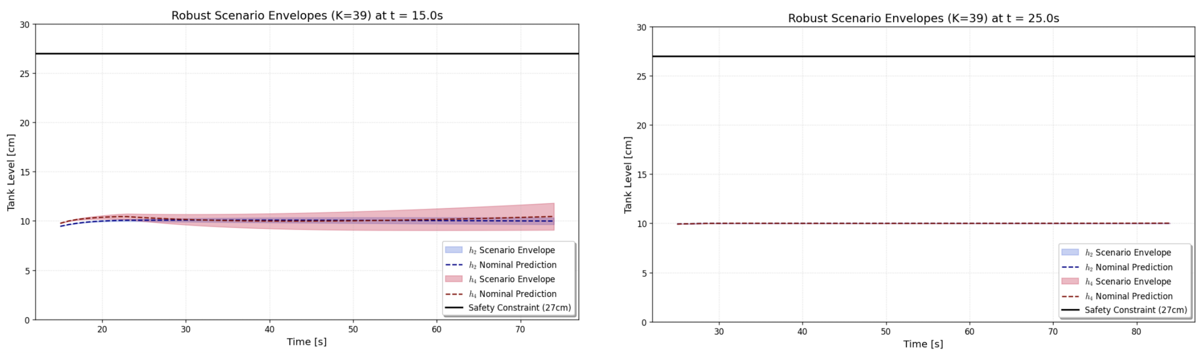

(a) SC-MPC envelops at 15s.

(b) SC-MPC envelops at 25s.

Figure 4.2: Prediction envelope.

Figure 4.2 shows the scenario envelopes at two time instants during closed-loop operation. Each subplot displays the $N = 60$ step prediction horizon starting from the current time: panel (a) at $t = 15$ s covers the window $[15, 75]$ s and panel (b) at $t = 25$ s covers $[25, 85]$ s. The shaded bands represent the spread of the $K = 39$ scenario trajectories around the nominal prediction (dashed lines) for both h_2 and h_4.

At $t = 15$ s the system is still in the early transient phase, so the scenario spread is wide: uncertain valve coefficients produce significantly different outflow predictions when the levels are far from steady-state. By $t = 25$ s the controlled outputs have nearly converged to the $[10, 10]$ cm target and the envelope has narrowed considerably. This contraction of uncertainties occurs because, near equilibrium, the level-dependent Torricelli derivatives $a_i^{(j)}$ vary less across scenarios, so all K realizations predict similar dynamics when the tanks are close to the setpoint. The safety constraint $h_{max} = 27$ cm remains well above the output uncertainty envelope in both cases, confirming that constraint satisfaction is maintained throughout the transient.

Table 4.3 summarises the SC-MPC parameters. All MPC tuning parameters (horizons, weights, input bounds) are identical to those in Table 4.2.

Table 4.3: SC-MPC parameters.

Parameter	**Symbol**	**Value**
Number of scenarios	K	39
Risk level	ϵ	5%
Support rank bound	ρ	$n_u = 2$
Valve uncertainty range	—	$[0.75,\, 0.90] \cdot A_i^{\mathrm{id}}\ for\ i = 1, ...4$
Slack penalty	α	10^6
RMSE safety margin	—	2.09 cm
Effective level bounds	$[h_{\min},\, h_{\max}]$	$[2.19,\, 24.91]$ cm

The computational overhead compared to successive linearization MPC comes primarily from the K-fold propagation of the free response in the tightening step and the K-fold trajectory propagation for the scenario envelope computation. Both operations involve matrix-vector products with the 6×6 augmented system matrix and scale linearly with K and N. In practice, all computations complete within the $T_s = 1$ s sampling period.

The scenario framework provides a principled trade-off between robustness and conservatism. The RMSE-based additional tightening adds a safety margin that accounts for modelling errors beyond the valve coefficient uncertainty; in practice this margin proved necessary to avoid constraint violations during experiments where the linear predictions deviated from the true nonlinear plant response. The framework is extended in Section 4.4 through the integration of online set membership identification, which allows the uncertainty set to adapt as new measurements become available.

The theoretical guarantee established in Section 2.4 is confirmed by the experimental results of Chapter 5. Across all conducted tests: setpoint tracking, multi-step references, large, signal operation near the actuator limit and disturbance rejection there were no violations of the effective level bounds $[h_{\min},\, h_{\max}] = [2.19,\, 24.91]$ cm were recorded for either h_2 or h_4. Given that these tests deliberately stress the controller with operating conditions where the linear prediction model is least accurate (large transients, coupling-induced cross-disturbances and near-saturation inputs), the absence of violations across the full experimental campaign provides strong empirical evidence that the true average violation rate remains well below the 5% ceiling.

4.4. Adaptive Scenario-Based MPC with Set Membership Identification

This section presents the most advanced controller in this thesis: an adaptive SC-MPC that integrates online SMI of the valve parameters with the scenario-based constraint enforcement framework of Section 4.3. While the SC-MPC samples from a fixed uncertainty set determined offline, the adaptive variant updates this set in real time using closed-loop measurements. The uncertainty set contracts as the controller learns from operating data, progressively reducing conservatism and improving tracking performance.

The approach builds on the framework of Tanaskovic et al. [10] and its extension to scenario-based formulations, which combines set membership estimation with scenario optimization. The key idea of this proposed SMI approach is to represent parametric uncertainty as an axis-aligned box in the space of valve outlet areas. In particular, this box is updated at each sample using the residual between the observed and predicted outflow for each tank, and then sample scenarios from the updated box to tighten the MPC output constraints. The resulting controller maintains the probabilistic constraint satisfaction guarantees of SC-MPC while adapting to the true system parameters over time.

The uncertain parameters are the four outlet valve coefficients $\theta = [A_1 \;\; A_2 \;\; A_3 \;\; A_4]^\top$, identical to the parameter vector used in the SC-MPC. The SMI framework assumes that these parameters lie within a compact set Θ_t that is refined at each time step t using new measurement data [11]. In the classical formulation, the measurement model takes the form

$$y_t = \varphi_t^\top \theta + e_t, \qquad |e_t| \leq \varepsilon, \tag{4.37}$$

where φ_t is a known regressor, e_t is the modelling error and ε is a known noise bound. At each time step a new strip (the set of parameters consistent with the latest measurement) is intersected with the current feasible parameter set to produce a refined set $\Theta_{t+1} \subseteq \Theta_t$ [10].

For the quadruple tank system, the SMI is performed directly on the physical mass balance. Each tank obeys

$$S \frac{dh_i(t)}{dt} = q_{\text{in},i}(t) - A_i \sqrt{2g\, h_i(t)}, \tag{4.38}$$

where A_i is time-invariant (the unknown constant to be identified) while $q_{\text{in},i}(t)$ and $h_i(t)$ vary with time. In the discrete-time implementation with sampling period T_s, t is replaced by the sample index k. Meanwhile, $q_{\text{in},i}$ is the total inflow to tank i and A_i is the uncertain outlet valve coefficient. Rearranging for A_i and using a backward Euler approximation for the derivative, $\dot{h}_i \approx \dfrac{h_i(k) - h_i(k-1)}{T_s}$, yields the observed value at sample k,

$$A_i^{\text{obs}}(k) = \frac{q_{\text{in},i}(k) - S\,\dfrac{h_i(k) - h_i(k-1)}{T_s}}{\sqrt{2g\,h_i(k-1)}}. \tag{4.39}$$

The denominator uses $h_i(k-1)$ because the outflow $A_i\sqrt{2g\,h_i}$ acting during the interval $[k{-}1, k)$ is driven by the level at the start of that interval. The inflow $q_{\text{in},i}$ is computed from the pump flow gains $\kappa_j(\cdot)$ and split ratios $\gamma_j(\cdot)$ evaluated at the previously applied voltage u_{k-1} using the identified lookup tables. These tables store experimentally identified values of κ_j and γ_j at discrete voltage points over $[4,\,18]\,\text{V}$ with a resolution of $0.1\,\text{V}$; values at intermediate voltages are recovered by linear interpolation and the input is clamped to $[4,\,18]\,\text{V}$ before querying to remain within the valid identification range. Specifically, the inflows are

$$q_{\text{in}}(k) = \begin{bmatrix} q_{\text{in},1}(k) \\ q_{\text{in},2}(k) \\ q_{\text{in},3}(k) \\ q_{\text{in},4}(k) \end{bmatrix} = \begin{bmatrix} \kappa_2(u_{k-1})\,\gamma_2(u_{k-1}) \\ \kappa_1(u_{k-1})\,(1-\gamma_1(u_{k-1})) \\ \kappa_1(u_{k-1})\,\gamma_1(u_{k-1}) \\ \kappa_2(u_{k-1})\,(1-\gamma_2(u_{k-1})) \end{bmatrix}, \tag{4.40}$$

consistent with the flow routing of (3.3). The quantity $\varphi_i(k) = \sqrt{2g\,h_i(k)}$ serves as the regressor for tank i because rearranging (4.38) gives

$$A_i \cdot \varphi_i(k) \approx q_{\text{in},i}(k) - S\,\frac{h_i(k) - h_i(k-1)}{T_s},$$

which is linear in the unknown parameter A_i with known coefficient $\varphi_i(k{-}1) = \sqrt{2g\,h_i(k-1)}$. Rearranging into the regression form of (4.37) gives

$$\underbrace{q_{\text{in},i}(k) - S\,\frac{h_i(k) - h_i(k-1)}{T_s}}_{y_k} = \underbrace{\sqrt{2g\,h_i(k-1)}}_{\varphi_{k-1}} \cdot A_i + e_k, \tag{4.41}$$

where the left-hand side is the measured net outflow, φ_{k-1} is the regressor evaluated at the previous sample (consistent with the denominator of (4.39)), and e_k subsumes linearization and measurement errors. Solving for A_i recovers (4.39). The strip at sample k for parameter i is

$$\mathcal{S}_{i,t} = \{A_i : |A_i - A_i^{\text{obs}}(k)| \le \varepsilon_{\text{smi}}\},$$

where $\varepsilon_{\text{smi}} = 0.35\,\text{cm}$ is the assumed noise bound, chosen approximately twice the sensor noise level of $\approx 0.15\,\text{cm}$ to provide robustness against transient estimation errors without being so large that the strip fails to contract the box.

The feasible parameter set is represented as an axis-aligned uncertainty box (interval hull). Since $\theta = [A_1,\ A_2,\ A_3,\ A_4]^\top$, the box for each component i at each time step k is an interval $[A^{(k)}_{\min,i},\ A^{(k)}_{\max,i}]$. Therefore, the full parameter box is defined as:

$$\Theta_k = [A^{(k)}_{\min,1},\ A^{(k)}_{\max,1}] \times \cdots \times [A^{(k)}_{\min,4},\ A^{(k)}_{\max,4}], \tag{4.42}$$

which is the simplest form of bounding-box set membership [10]. The box representation was chosen because the parameters appear independently in the tank dynamics (each A_i affects only tank i) and the box update is computationally trivial and compatible with real-time control operation.

The box update at each eligible time step k follows a two-phase procedure:

1. *Expansion phase:* the current box is slightly widened to prevent collapse and preserve the ability to track slow parameter drift.
2. *Contraction phase:* a strip intersection with the latest measurement tightens the box around the true parameter.

In the first phase, a Minkowski expansion with factor $\lambda_{\text{exp}} = 0.005$ is applied once every 60 samples (60 s at $T_s = 1$ s),

$$A_{\min,i} \leftarrow A_{\min,i} - \lambda_{\text{exp}}(A_{\max,i} - A_{\min,i}), \qquad A_{\max,i} \leftarrow A_{\max,i} + \lambda_{\text{exp}}(A_{\max,i} - A_{\min,i}). \tag{4.43}$$

Note that without expansion the box can become too tight to include the true parameter after a noise spike, causing the SMI to lose the true system. The expansion is applied infrequently (every 60 steps) rather than at every timestep because, for example, applying it at every step with any nonzero factor would systematically counteract the contraction from strip intersections, preventing the box from shrinking. The infrequent application with $\lambda_{\text{exp}} = 0.005$ provides the necessary collapse prevention while allowing genuine shrinkage to accumulate between expansion events. After the expansion step, each component is immediately clamped to the initial uncertainty range $[A^{(0)}_{\min,i},\ A^{(0)}_{\max,i}] = [0.75\, A^{\text{id}}_i,\ 0.90\, A^{\text{id}}_i]$,

$$A_{\min,i} \leftarrow \max\bigl(A_{\min,i},\ A^{(0)}_{\min,i}\bigr), \qquad A_{\max,i} \leftarrow \min\bigl(A_{\max,i},\ A^{(0)}_{\max,i}\bigr), \tag{4.44}$$

ensuring that the expansion cannot push the box beyond the physically motivated initial range.

In the second phase, applied at every eligible discrete step k independently of the expansion schedule, the strip intersection contracts the box. For each component i, a candidate interval is first computed,

$$A^{\mathrm{cand}}_{\min,i} \leftarrow \max\bigl(A_{\min,i},\ A^{\mathrm{obs}}_i(k) - \varepsilon_{\mathrm{smi}}\bigr), \qquad A^{\mathrm{cand}}_{\max,i} \leftarrow \min\bigl(A_{\max,i},\ A^{\mathrm{obs}}_i(k) + \varepsilon_{\mathrm{smi}}\bigr), \quad (4.45)$$

yielding the following candidate box. The box is updated only if the candidate interval is non-degenerate (see below). Since this contraction occurs at every eligible step, genuine learning accumulates continuously between the infrequent expansion events. No hard clamp is applied after the strip intersection, allowing the box to shrink below the initial range as measurements confirm the true parameter location.

The finite-difference derivative in (4.39) amplifies measurement noise and is unreliable during transients. For that reason, several safety mechanisms are followed to prevent the SMI from corrupting the parameter box under adverse conditions:

- *Guard period.* First, the SMI update is activated only after $k > 40$ discrete time steps, allowing the system to settle from its initial condition toward the target setpoint. During this initial transient, the levels change rapidly and the finite-difference derivative produces unreliable estimates of A^{obs}_i.

- *Reference-change cooldown.* Then, at each discrete time step k the controller monitors the reference signal y_{ref}. If a reference change is detected ($\|y_{\mathrm{ref},k} - y_{\mathrm{ref},k-1}\|_\infty > 0.01$), a cooldown counter is set to $n_{\mathrm{cool}} = 100$ discrete time steps and decremented at each subsequent step. The SMI update is inhibited whenever the cooldown counter is positive. This mechanism prevents learning during the transient that follows a setpoint change, where the mass balance inversion (4.39) would produce biased estimates due to the large and rapidly varying derivative.

  Even outside the cooldown period, the strip intersection (4.45) is applied only when the filtered derivative satisfies $|\dot{h}_i| < 0.3\,\mathrm{cm/s}$. This provides a per-tank guard against residual transients that the cooldown may not fully capture.

- *Regressor threshold.* The update for tank i is skipped when the outflow regressor $\varphi_i(k) = \sqrt{2g\,h_i(k)}$ falls below 0.5, preventing updates from unreliable data when tank levels are near zero.

- *Rejection of degenerate intersections.* If the candidate interval is degenerate ($A^{\mathrm{cand}}_{\max,i} - A^{\mathrm{cand}}_{\min,i} < 10^{-4}$), the intersection is rejected and the current box is retained unchanged. Resetting to the initial range would discard all previously accumulated learning, so the conservative choice is to preserve the current box rather than revert to a wider but less informed set.

The raw level measurements are filtered using a first-order exponential moving average,

$$\hat{x}_k = \alpha_f\, \hat{x}_{k-1} + (1 - \alpha_f)\, x_k, \qquad \alpha_f = 0.5, \tag{4.46}$$

before using the finite-difference derivative in (4.39) to compute A_i^{obs}. This filtering is essential because the derivative $\dot{h}_i$ amplifies measurement noise and unfiltered derivatives would produce excessively noisy estimates of A_i^{obs}.

The complete SMI procedure including all safety mechanisms is summarized in Algorithm 1.

At each time step k, $K = 39$ scenarios are drawn from the current parameter box Θ_k, preserving the 5% risk level guarantee. The sampling procedure generates each scenario as

$$\theta_k^{(j)} = \theta_{\text{center}} + \theta_{\text{radius}} \odot \xi^{(j)}, \qquad \xi^{(j)} \sim \mathcal{U}([-1, 1]^4), \quad j = 1, \ldots, K, \tag{4.47}$$

where $\theta_{\text{center}} = (\theta_{\min} + \theta_{\max})/2$ and $\theta_{\text{radius}} = (\theta_{\max} - \theta_{\min})/2$ are updated at every sample from the current SMI box.

Recall that, the scenarios are resampled independently at each time step, ensuring that the constraint tightening is not biased by a particular fixed draw. As the SMI box contracts over time, the center converges toward the true parameter and the radius decreases, producing scenarios that cluster more tightly around the true system. This contraction progressively reduces the conservatism of the scenario-tightened constraints, improving closed-loop performance.

The nominal prediction model is constructed using successive linearization at the current operating point, following the same procedure as Subsection 4.2.2. The continuous-time Jacobians $A_c(x_k)$ and $B_c(u_{k-1})$ follow equations (4.20)–(4.21), discretized via Forward Euler. The nominal valve parameters used for computing the Jacobian are the identified values A^{id} (not the learned box center θ_{center}), ensuring that the prediction model remains anchored to the best available plant identification regardless of the current state of the SMI. The learned box affects only the scenario generation and constraint tightening, not the nominal prediction.

The velocity-form augmentation and prediction matrices S_x, S_u are rebuilt online exactly as in the successive linearization MPC.

The constraint tightening follows the same procedure as for SC-MPC. For each set of $K = 39$ scenarios, a scenario-specific augmented matrix $A_{\text{aug}}^{(j)}$ is constructed using the

Algorithm 1 Set membership identification with safety mechanisms.

Require: Measurement x_k, previous filtered state $\hat{x}_{k-1}$, inflow $q_{\text{in},k}$, current box $[A^k_{\min,i}, A^k_{\max,i}]$ for all $i = 1, \ldots, 4$, reference $y_{\text{ref},k}$, previous output reference $y_{\text{ref},k-1}$, cooldown counter c, sample count k, guard counter r

Ensure: Updated box $[A^{k+1}_{\min,i}, A^{k+1}_{\max,i}]$ for all i, updated filtered state $\hat{x}_k$, updated cooldown c

$\hat{x}_k \leftarrow 0.5\,\hat{x}_{k-1} + 0.5\,x_k$
if $\|y_{\text{ref},k} - y_{\text{ref},k-1}\|_\infty > 0.01$ **then**
  $c \leftarrow 100$ ▷ Trigger cooldown
  $r \leftarrow 0$ ▷ Reset expansion guard
end if
$c \leftarrow \max(c-1,\ 0)$
if $k > 40$ **and** $c = 0$ **then**
  $r \leftarrow r + 1$
  if $r \bmod 60 = 0$ **then** ▷ Expand only every 60 time steps
    Expand box via (4.43)
    Clamp via (4.44) ▷ Clamp only after expansion
  end if
  for $i = 1, \ldots, 4$ **do** ▷ Strip intersection at every eligible step
    $\varphi_i \leftarrow \sqrt{2g\ \max(h_{i,k-1},\ 0.1)}$
    if $\varphi_i > 0.5$ **then**
      $\dot{h}_i \leftarrow (\hat{x}_{k,i} - \hat{x}_{k-1,i})/T_s$
      if $|\dot{h}_i| < 0.3$ **then**
        Compute A^{obs}_i via (4.39)
        $A^{\text{cand}}_{\min,i} \leftarrow \max\left(A_{\min,i},\ A^{\text{obs}}_i - \varepsilon_{\text{smi}}\right)$
        $A^{\text{cand}}_{\max,i} \leftarrow \min\left(A_{\max,i},\ A^{\text{obs}}_i + \varepsilon_{\text{smi}}\right)$
        if $A^{\text{cand}}_{\max,i} - A^{\text{cand}}_{\min,i} > 10^{-4}$ **then**
          $A_{\min,i} \leftarrow A^{\text{cand}}_{\min,i}, \quad A_{\max,i} \leftarrow A^{\text{cand}}_{\max,i}$
        else
          Keep current box $[A_{\min,i},\ A_{\max,i}]$
        end if
      end if
    end if
  end for ▷ No clamp here: allow genuine shrinkage below initial range
end if

sampled valve coefficients $\theta_k^{(j)}$ from (4.47). The free response is propagated for $N = 60$ steps and the worst-case deviations across all scenarios are used to tighten the output bounds, with $h_{\min} = 2.19\,\mathrm{cm}$ and $h_{\max} = 24.91\,\mathrm{cm}$ as in Section 4.3.

The critical difference from the non-adaptive SC-MPC is that, in this case, the scenario spread decreases over time as the SMI box contracts. In the initial phase, when the box is wide ($[\theta_{\min}, \theta_{\max}] = [0.75\,A^{\mathrm{id}},\ 0.90\,A^{\mathrm{id}}]$), the tightened constraints are conservative, identical to the fixed SC-MPC. As measurements accumulate and the box shrinks, the scenario spread narrows and the tightened bounds approach the nominal (untightened) bounds, recovering performance close to the successive linearization MPC. This adaptation mechanism represents the practical advantage of the adaptive scheme: it automatically transitions from robust to nominal behaviour as the model uncertainty is decreased.

The QP formulation of the adaptive SC-MPC is identical to the SC-MPC. The decision variable is $[\Delta U^\top\ \sigma]^\top$ with $\sigma \geq 0$ penalized by $\alpha = 10^6$. The constraint matrix (4.36) and slack relaxation (4.35) are assembled identically. The QP problem is solved using `quadprog` with the active-set algorithm and warm-started from the full previous solution vector $[\Delta U_{k-1}^\top\ \sigma_{k-1}]^\top$. If the solver returns a non-positive exit flag, the control increment is set to $\Delta u_k = 0$, holding the previous command u_{k-1}.

After the QP is solved, the same $K = 39$ scenarios drawn at the current step are propagated forward using the optimal control sequence $\Delta U^\star$ to compute the scenario envelope, following the same procedure as for SC-MPC.

Table 4.4 summarises the adaptive SC-MPC parameters. All MPC tuning parameters (horizons, weights, input bounds) are identical to those in Table 4.2.

Table 4.4: Adaptive SC-MPC parameters.

Parameter	**Symbol**	**Value**
Number of scenarios	K	39
Risk level	ϵ	5%
Initial uncertainty range	Θ_0	$[0.75,\ 0.90] \cdot A^{\mathrm{id}}$
Slack penalty	α	10^6
RMSE safety margin	δ_{RMSE}	2.09 cm
Effective level bounds	$[h_{\min},\ h_{\max}]$	[2.19, 24.91] cm
SMI parameters		
Noise bound	$\varepsilon_{\mathrm{smi}}$	0.35 cm
Expansion factor	$\lambda_{\exp}$	0.005 (applied every 60 steps)
Guard period	k_{guard}	40
Reference-change cooldown	n_{cool}	100
Derivative threshold	$\dot{h}_{\max}$	0.3 cm/s
Regressor threshold	$\varphi_{\min}$	0.5
EMA filter coefficient	α_f	0.5

5 | Performance Metrics and Experimental Results

This chapter evaluates and compares the implemented control strategies on the physical quadruple-tank plant. The analysis is organised by experimental scenarios that stress different closed-loop properties: setpoint tracking, multi-step tracking under symmetric and asymmetric references, large-signal operation near actuator limits, feasibility under asymmetric steady-state targets and disturbance rejection under unmeasured manual water injections. Unless stated otherwise, the controlled outputs are the water levels of Tanks 2 and 4 (h_2 and h_4), while Tanks 1 and 3 remain uncontrolled.

5.1. Performance Evaluation Criteria

To quantitatively evaluate the performance of the controllers the following metrics are employed.

1. *Integral Absolute Error (IAE):* Measures cumulative tracking deviation over the entire experiment.

$$\mathrm{IAE} = \sum_{k=0}^{N} |r(k) - y(k)| \tag{5.1}$$

2. *Root Mean Square Error (RMSE):* Provides an average measure of error magnitude that penalises large deviations.

$$\mathrm{RMSE} = \sqrt{\frac{1}{N}\sum_{k=1}^{N}(r(k) - y(k))^2} \tag{5.2}$$

3. *Steady-State Error (e_{ss}):* The mean and maximum absolute tracking error computed over the last 50 s of the experiment. This window captures the regulation performance once transients have decayed.

4. *Settling Time (τ_{set}):* The earliest time after which the output enters and remains within a 5% band of the step size around the final reference value for at least 10 consecutive seconds, where $\Delta y = y_\infty - y_0$ denotes the reference step amplitude.

$$\tau_{\mathrm{set}} = \inf\{t > 0 : |y(\tau) - y_\infty| \leq 0.05\,|\Delta y|,\ \forall\,\tau \in [t,\, t+10]\} \tag{5.3}$$

5. *Percentage Overshoot (OS):* The peak excursion beyond the reference normalised by the step size.

$$\mathrm{OS}(\%) = \frac{y_{\max} - y_\infty}{|y_\infty - y_0|} \times 100 \tag{5.4}$$

6. *Total Control Energy (E):* The sum of squared pump voltages over the experiment. Both pumps are combined because the cross-coupled plant structure means that each pump affects both controlled outputs.

$$E = \sum_{k=0}^{N} \left[u_1^2(k) + u_2^2(k)\right] \tag{5.5}$$

7. *Total Variation (TV):* The cumulative sum of absolute input increments across both pumps. This metric quantifies actuator activity and is directly related to mechanical wear.

$$\mathrm{TV} = \sum_{k=1}^{N} |u_1(k) - u_1(k{-}1)| + \sum_{k=1}^{N} |u_2(k) - u_2(k{-}1)| \tag{5.6}$$

5.2. Test 1: Setpoint Tracking

This test evaluates the setpoint tracking performance of five control strategies: LQI, Nominal MPC, SL-MPC, SC-MPC with open-loop startup and SC-MPC without open-loop startup. All experiments share a common duration of 120 s and a data logging rate of 0.01 s. The first four controllers employ a 15 s open-loop startup phase during which the pumps are driven at the equilibrium voltages corresponding to the target reference of $h_2 = h_4 = 10$ cm. In this phase the integrator is disabled and the controller output is held constant. The fifth configuration launches the SC-MPC controller from the beginning without any open-loop prefill, providing a direct comparison of startup behaviour under identical initial conditions.

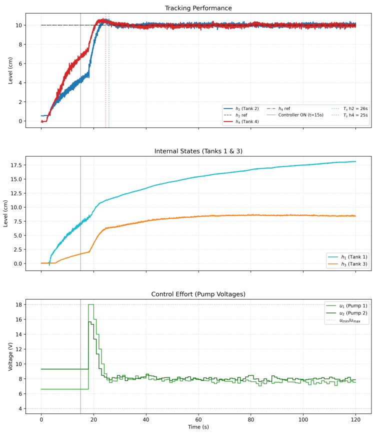

Figure 5.1: LQI

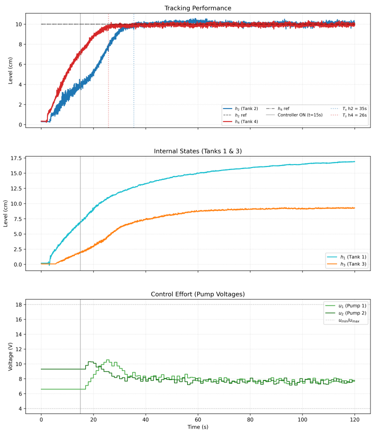

Figure 5.2: Nominal MPC

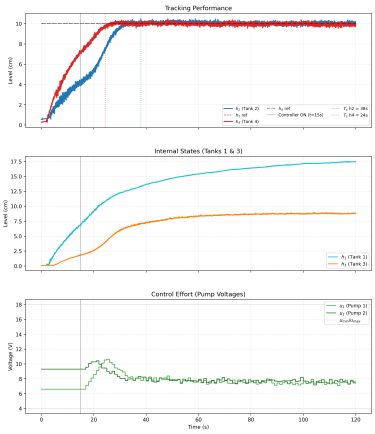

Figure 5.3: SL-MPC

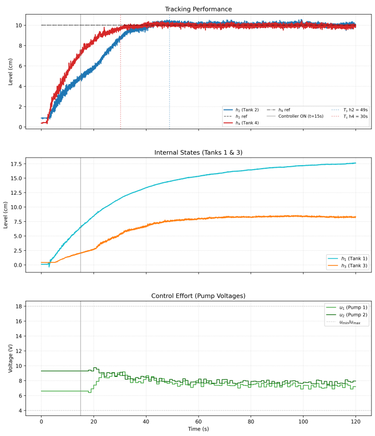

Figure 5.4: Open-loop SC-MPC

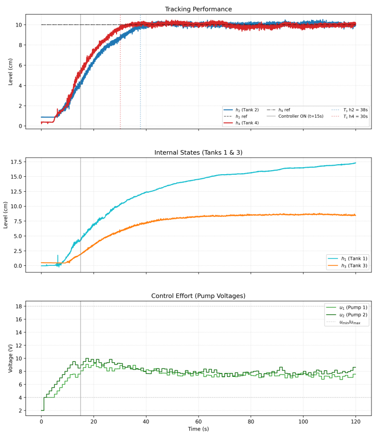

Figure 5.5: Closed-loop SC-MPC

The tracking and control effort metrics for each controller are summarised in Table 5.1. All values are computed over the full 120 s experiment window. The column labelled SC-MPC* denotes the variant without the open-loop startup phase.

Table 5.1: Performance metrics comparison for Test 1: Setpoint Tracking ($h_2 = h_4 = 10\,\mathrm{cm}$).

Metric	Tank	LQI	Nominal MPC	SL-MPC	SC-MPC	SC-MPC*
IAE (cm s)	2	148.4	184.8	173.3	174.3	181.5
	4	121.3	117.8	114.7	117.9	164.5
RMSE (cm)	2	2.910	3.204	3.053	2.893	3.143
	4	2.542	2.421	2.383	2.345	3.041
e_{ss} mean (cm)	2	0.0768	0.0822	0.0794	0.1103	0.0977
	4	0.0748	0.0863	0.0818	0.0922	0.1356
e_{ss} max (cm)	2	0.3804	0.4947	0.4131	0.5567	0.5175
	4	0.4319	0.4907	0.4793	0.4923	0.6442
Settling Time (s)	2	25.8	35.5	38.0	49.0	37.8
	4	24.6	25.8	24.4	30.3	30.2
Overshoot (%)	2	6.85	5.85	5.47	6.10	5.67
	4	6.94	3.67	3.94	4.43	4.77
Total Energy (V^2 s)	—	16 568	15 507	15 245	14 954	14 507
Total Variation (V)	—	71.4	64.8	56.2	62.7	81.0

The LQI controller delivers the fastest transient response with settling times of 25.8 s and 24.6 s for Tanks 2 and 4 respectively. It achieves the lowest IAE for Tank 2 (148.4 cm s) and the smallest steady-state errors across both tanks. This aggressive tracking comes at the cost of the highest control energy and the largest input variation among the four open-loop startup configurations (71.4 V). The pump 1 voltage saturates at 18 V immediately after the controller engages at $t = 15\,\mathrm{s}$.

Nominal MPC reduces total control energy to 15507 V^2s compared to LQI while maintaining comparable steady-state accuracy. The settling time for Tank 2 increases to 35.5 s because the MPC cost function penalises aggressive input changes. Overshoot is notably lower for Tank 4 (3.67% versus 6.94% for LQI) and the control signal remains within the saturation limits throughout the experiment.

SL-MPC achieves the lowest total input variation (56.2 V) among all five configurations and the lowest IAE for Tank 4 (114.7 cm s). The online re-linearisation of the prediction model at each sampling instant produces smoother control trajectories without sacrificing steady-state accuracy. The settling time for Tank 2 (38.0 s) is slightly longer than Nominal MPC because the model updates introduce small transient adjustments as the operating point evolves during the fill phase.

SC-MPC achieves the lowest total control energy (14954 V^2s) among the four open-loop startup controllers. The constraint tightening introduced by the scenario-based robustification produces a more conservative transient response with the longest settling time for Tank 2 (49.0 s). The steady-state error is slightly larger than the other MPC variants ($e_{ss,\text{mean}} = 0.110$ cm for Tank 2) due to the tightened output constraints that prevent the controller from utilising the full feasible region.

Removing the open-loop prefill phase reveals several differences. The settling time for Tank 2 improves from 49.0 s to 37.8 s because the controller drives the system from the start rather than relying on constant equilibrium voltages during the first 15 s. However the Tank 4 IAE increases to 164.5 cm s (compared to 117.9 cm s with startup) because the initial state is far from the linearisation point and the robust constraint tightening limits the available control authority. The total input variation rises to 81.0 V, the highest among all configurations, reflecting the more aggressive actuation required when starting from near-empty tanks without a preliminary fill. Control energy remains the lowest overall at 14507 V^2s.

The results reveal a consistent trade-off between transient speed and control economy. LQI provides the fastest settling but consumes the most energy and produces the highest actuator activity among the open-loop startup configurations. The three MPC variants with open-loop startup occupy a middle ground: Nominal MPC and SL-MPC offer similar tracking performance with SL-MPC distinguishing itself through the smoothest control signal. SC-MPC sacrifices settling speed in exchange for the lowest energy consumption and built-in robustness to parametric uncertainty.

The comparison between SC-MPC with and without the open-loop startup phase highlights the importance of the initialisation strategy. The open-loop prefill brings the system close to the operating point before the controller engages, which benefits the linearisation-based prediction model and avoids large constraint tightening penalties during the initial transient. Without this phase the controller must handle the nonlinear fill dynamics from near-empty conditions, resulting in higher actuator wear despite faster Tank 2 settling.

5.3. Test 2: Step Reference Tracking

Test 2 evaluates the four controllers on a multi-step reference trajectory applied simultaneously to both controlled tanks (h_2 and h_4). The reference sequence is $0 \to 10 \to 12 \to 10$ cm, producing three distinct phases: a large initial step-up from empty, a smaller 2 cm step-up and a 2 cm step-down. This scenario tests initial transient behaviour, incremental tracking precision and the ability to reject the coupling disturbance introduced by the cross-flow between tanks. The per-phase settling times and overshoot values in Table 5.3

complement the global metrics of Table 5.2 by isolating how each controller handles the large filling transient separately from the smaller incremental steps.

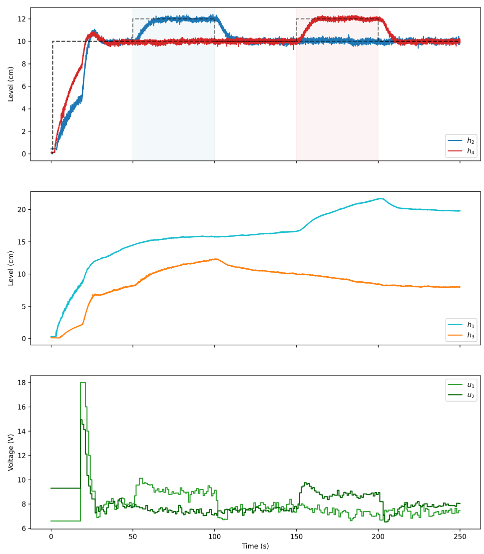

Figure 5.6: LQI

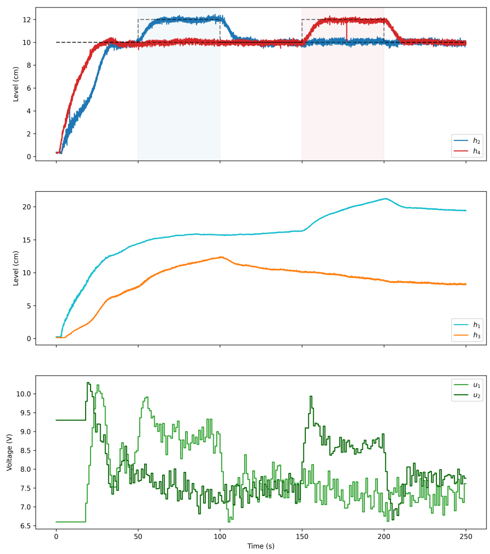

Figure 5.7: Nominal MPC

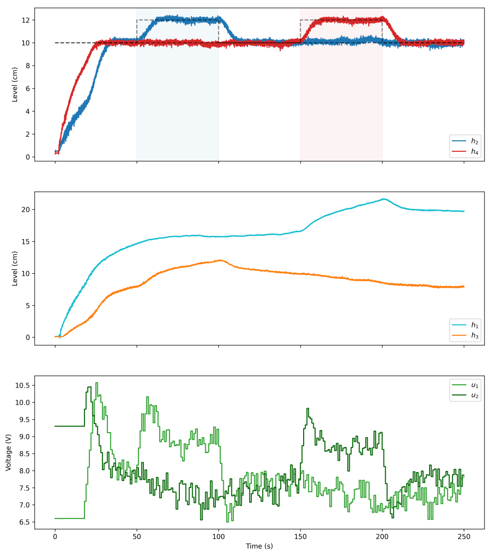

Figure 5.8: SL-MPC

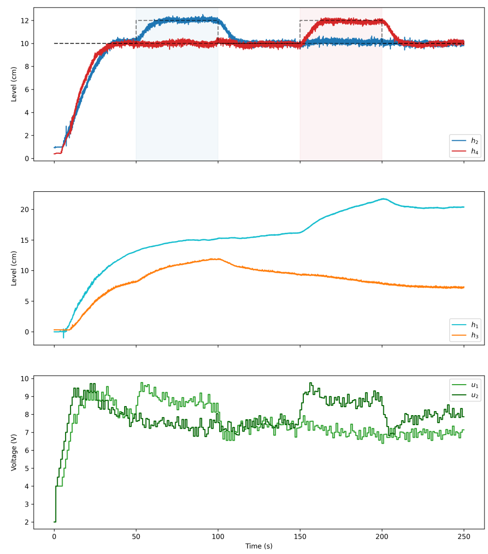

Figure 5.9: Closed-loop SC-MPC

LQI achieves the lowest combined IAE across both tanks (77.0 and 56.7 cm s for h_2 and h_4 respectively) and the lowest RMSE on both channels. During the initial $0 \rightarrow 10$ cm phase, LQI settles in 29.6 s on h_2 and 29.8 s on h_4, the fastest among all four controllers. On the incremental $10 \to 12$ cm and $12 \to 10$ cm steps, settling times remain short (12.0 s and 11.7 s on h_2) but overshoot is consistently the highest of all controllers, reaching 24.00% on the down-step of h_2. This behaviour is consistent with Test 1: aggressive integral action drives fast convergence at the expense of pump voltage saturation at 18 V and elevated actuator activity (126.8 V total variation). The total energy of 31 281 V^2 s is the highest in this test, confirming that the speed advantage is purchased through greater actuation effort.

Nominal MPC is the slowest controller during the initial $0 \to 10$ cm phase on h_2, with a settling time of 39.8 s, reflecting the limitation of a fixed linear prediction model operating far from the nominal linearisation point during large transients. Performance on h_4 is more competitive at 26.3 s, since the level-dependent outflow nonlinearity affects h_2 more severely when filling from empty. On the incremental steps, Nominal MPC achieves competitive settling times (10.0 s on h_2 for the up-step) but produces the largest overshoot on the $12 \to 10$ cm return (26.04% on h_2), indicating sensitivity to model mismatch when the operating point shifts away from the nominal linearisation region. Steady-state accuracy is good, with mean errors of 0.0911 and 0.0833 cm on h_2 and h_4. Control energy (29 747 V^2 s) and total variation (131.5 V) are mid-range.

SL-MPC delivers the best performance on h_4 across all phases, achieving the lowest IAE (55.8 cm s) and the fastest settling times during both the initial fill (24.4 s) and the $10 \to 12$ cm up-step (8.9 s). This advantage is directly attributable to the online re-linearisation: as the h_4 level rises through the operating range, the updated Jacobian tracks the changing Torricelli gain and maintains prediction accuracy throughout the transient. On h_2, however, SL-MPC shows the highest IAE of all controllers (106.6 cm s), driven by a slower initial fill (32.0 s settling). This asymmetry arises because the cross-coupling effect on h_2 is captured less accurately during simultaneous filling from empty: the re-linearisation at the current h_2 level does not fully compensate for the coupled inflow dynamics when both tanks start from zero. SL-MPC achieves the lowest total variation (125.5 V), confirming the smoothest actuator behaviour observed in Test 1, and overshoot on the incremental steps is moderate (17.80% on the h_2 down-step), lower than both LQI and Nominal MPC on this transition.

SC-MPC produces the lowest total control energy (29 118 V^2 s) in this test, consistent with its energy-efficient behaviour in Test 1. The scenario-based constraint tightening delays the initial $0 \to 10$ cm settling to 33.9 s and 33.4 s on h_2 and h_4, as the robust output bounds limit the available control authority during the large filling transient. On the incremental up-step, SC-MPC exhibits the largest overshoot on h_2 (25.88%): when both channels simultaneously demand flow, the constraint margins are temporarily saturated, causing a larger transient excursion before the tightened bounds become binding again. On h_4, overshoot is the lowest among all controllers on the up-step (11.80%), reflecting the asymmetric effect of the scenario tightening across the two channels. Steady-state accuracy remains competitive (0.0813 cm mean error on h_2), while input variation (144.8 V) is the highest in this test, driven by more frequent corrective actions near the tightened bounds.

The results across all three phases reveal a complementary pattern of strengths. LQI dominates on speed and global IAE but consistently produces the highest overshoot and energy use. SL-MPC is the clear winner on h_4 tracking across all transitions and delivers the smoothest control signal, but its h_2 performance during simultaneous filling from empty is its weakest point. SC-MPC and Nominal MPC offer more balanced cross-channel behaviour, with SC-MPC trading settling speed for energy efficiency and Nominal MPC providing adequate performance at moderate cost.

A consistent observation across all four controllers is that the incremental $10 \to 12$ and $12 \to 10$ cm steps produce significantly higher overshoot values than the equivalent amplitude changes observed in Test 1. This is explained by the compressed experiment window: after the large $0 \to 10$ cm fill, the upper tank levels h_1 and h_3 are still rising toward their new equilibria when the next reference change is applied, so cross-coupling acts as an active disturbance during the incremental transitions. Controllers that do not explicitly account for this coupling effect, such as LQI and Nominal MPC, are more affected than SL-MPC, which partially compensates through its updated prediction model.

Table 5.2: Global metrics for Test 2 ($0 \to 10 \to 12 \to 10$ cm reference).

	IAE (cm·s)		RMSE (cm)		SS Error (cm)		Energy / Variance	
Controller	h_2	h_4	h_2	h_4	h_2	h_4	Total	Var
LQI	77.0	56.7	0.876	0.548	0.0831	0.0702	31 281	126.8
SC-MPC	90.2	79.0	0.886	0.742	0.0813	0.1045	29 118	144.8
SL-MPC	106.6	55.8	1.128	0.518	0.1095	0.0903	29 842	125.5
Nominal MPC	102.7	57.3	1.120	0.521	0.0911	0.0833	29 747	131.5

Table 5.3: Transition metrics for Test 2. T_s = settling time, OS = overshoot.

	$0 \to 10$ cm		$10 \to 12$ cm (Up)		$12 \to 10$ cm (Down)	
Controller	$T_s^{h_2}$ (s)	$T_s^{h_4}$ (s)	T_s / OS% (h_2)	T_s / OS% (h_4)	T_s / OS% (h_2)	T_s / OS% (h_4)
LQI	29.6	29.8	12.0 / 19.84	9.9 / 14.73	11.7 / 24.00	9.7 / 21.51
SC-MPC	33.9	33.4	11.2 / 25.88	11.6 / 11.80	11.1 / 20.74	10.5 / 22.49
SL-MPC	32.0	24.4	11.0 / 22.45	8.9 / 16.53	11.2 / 17.80	10.9 / 20.45
Nominal MPC	39.8	26.3	10.0 / 22.45	10.8 / 15.63	12.2 / 26.04	10.1 / 19.06

Table 5.4: Control effort comparison for Test 2.

Controller	Total Energy	Input Variance	Voltage Range (V)
LQI	31 281	126.8	6–18
SC-MPC	29 118	144.8	2–10
SL-MPC	29 842	125.5	6.5–10.5
Nominal MPC	29 747	131.5	6.5–10

LQI achieves the lowest combined IAE and the tightest tracking, particularly on Tank 2. SL-MPC performs best on Tank 4, which is consistent with the benefit of updating the linearisation point as the operating level changes. SC-MPC reduces energy but shows a more conservative transient on Tank 2. Nominal MPC provides balanced behaviour but is the slowest during the initial $0 \to 10$ cm filling on Tank 2, reflecting the limitations of a fixed linear model far from the nominal operating point.

5.4. Test 3: Simultaneous Pulse Tracking

Test 3 applies the same $0 \to 10 \to 12 \to 10$ cm reference sequence to both controlled tanks but introduces a simultaneous pulse phase where both references step up and then return together. Compared to Test 2, the experiment window is shorter, compressing the transients and leaving less time for steady-state convergence. This scenario stresses the controllers' ability to handle concurrent large-signal excitations with strong cross-coupling, since both pumps must respond to identical demand changes at the same time. The per-phase settling times and overshoot values in Table 5.6 complement the global metrics of Table 5.5 by isolating how the simultaneous excitation affects each channel independently.

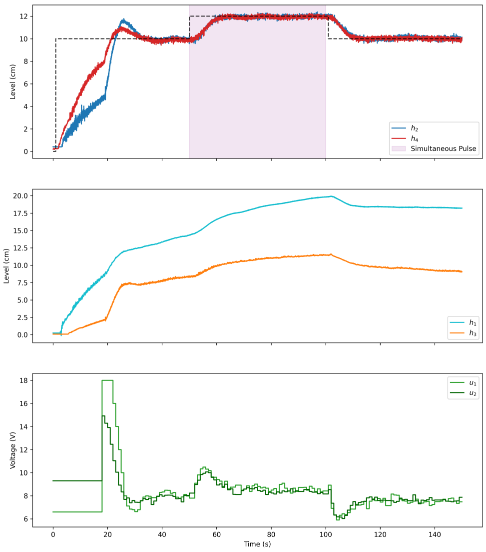

Figure 5.10: LQI

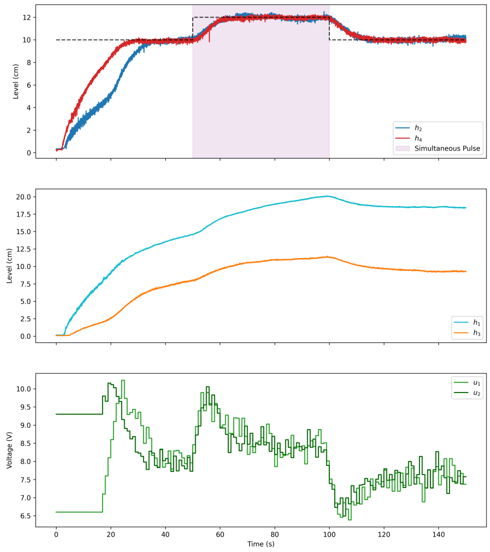

Figure 5.11: Nominal MPC

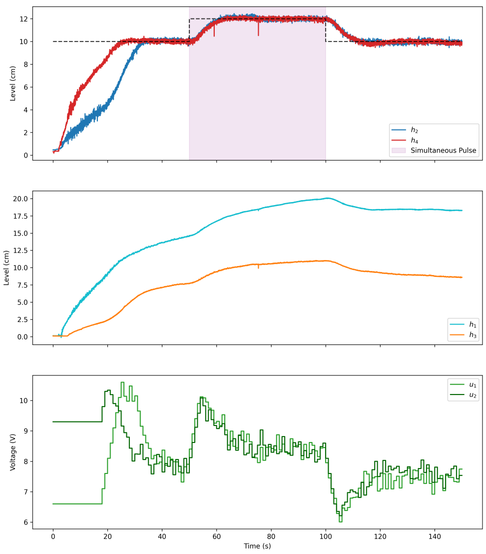

Figure 5.12: SL-MPC

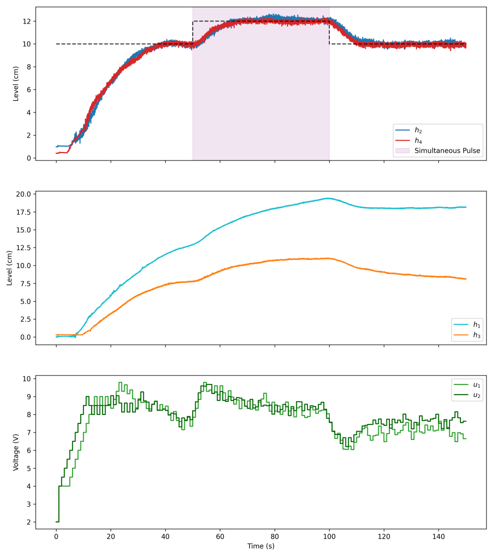

Figure 5.13: Closed-loop SC-MPC

Table 5.5: Global metrics for Test 3 simultaneous ($0 \to 10 \to 12 \to 10$ cm).

	IAE (cm·s)		RMSE (cm)		SS Error (cm)		Energy / Variance	
Controller	h_2	h_4	h_2	h_4	h_2	h_4	Total	Var
LQI	72.7	48.9	1.223	0.698	0.1010	0.0762	19 648	88.5
SC-MPC	84.0	77.6	1.167	1.115	0.0895	0.0966	17 607	84.8
SL-MPC	95.1	46.7	1.583	0.672	0.0950	0.0921	18 165	76.1
Nominal MPC	89.9	48.3	1.422	0.688	0.1019	0.0773	18 159	80.5

Table 5.6: Transition metrics for Test 3 simultaneous.

	$0 \to 10$ cm		$10 \to 12$ cm (Up)		$12 \to 10$ cm (Down)	
Controller	$T_s^{h_2}$ (s)	$T_s^{h_4}$ (s)	T_s / OS% (h_2)	T_s / OS% (h_4)	T_s / OS% (h_2)	T_s / OS% (h_4)
LQI	31.9	31.1	10.3 / 19.10	10.2 / 12.77	9.9 / 20.66	8.8 / 20.37
SC-MPC	36.3	36.5	12.6 / 31.84	13.7 / 13.10	14.1 / 23.27	9.0 / 29.76
SL-MPC	32.1	24.5	9.4 / 22.69	10.5 / 16.29	9.9 / 27.19	10.2 / 26.58
Nominal MPC	34.6	26.3	11.0 / 25.23	10.5 / 14.08	11.1 / 28.00	10.9 / 19.15

Table 5.7: Control effort comparison for Test 3 simultaneous.

Controller	Total Energy	Input Variance	Voltage Range (V)
LQI	19 648	88.5	6–18
SC-MPC	17 607	84.8	2–10
SL-MPC	18 165	76.1	6–10.5
Nominal MPC	18 159	80.5	6.5–10

LQI again achieves the lowest combined IAE across both tanks (72.7 and 48.9 cm s for h_2 and h_4 respectively) and the lowest RMSE on both channels. The initial $0 \to 10$ cm settling times are 31.9 s and 31.1 s, slightly slower than in Test 2, which is consistent with the increased cross-coupling load when both references change simultaneously. On the incremental steps, LQI maintains short settling times (10.3 s on h_2 for the up-step and 9.9 s for the down-step) while overshoot remains moderate relative to the other tests at 19.10% and 20.66% on h_2. The total energy of 19 648 V^2 s and total variation of 88.5 V are lower than in Test 2, reflecting the shorter experiment window rather than a reduction in actuation aggressiveness. Pump voltage still reaches 18 V during the initial fill, confirming the same saturation behaviour observed in previous tests.

Nominal MPC settles the initial $0 \rightarrow 10$ cm phase in 34.6 s on h_2 and 26.3 s on h_4, showing the same asymmetric channel behaviour as in Test 2. The fixed linear model struggles more with simultaneous excitation than with the sequential case: the h_2 RMSE rises to 1.422 cm compared to 1.120 cm in Test 2, driven by the increased cross-coupling load during the concurrent fill. On the incremental steps, settling times are competitive (11.0 s on h_2 for the up-step) but overshoot on the down-step reaches 28.00% on h_2, the highest recorded for Nominal MPC across all three phases of this test. Steady-state accuracy remains good, with mean errors of 0.1019 and 0.0773 cm on h_2 and h_4. Control energy (18 159 V^2 s) and total variation (80.5 V) are mid-range.

SL-MPC continues to deliver the best h_4 performance, achieving the lowest IAE on that channel (46.7 cm s) and the fastest initial settling on h_4 (24.5 s). The online re-linearisation maintains prediction accuracy on h_4 even under simultaneous excitation, confirming that the updated Jacobian effectively tracks the Torricelli gain variation across both channels when they move together. On h_2, however, SL-MPC again shows the highest IAE (95.1 cm s) and a slower initial settling (32.1 s), consistent with the coupling asymmetry observed in Test 2. SL-MPC achieves the lowest total variation (76.1 V) of all controllers in this test, preserving its characteristic smoothness even under the higher cross-coupling load of simultaneous excitation. Overshoot on the incremental steps is moderate, with 22.69% on the h_2 up-step and 27.19% on the down-step, both higher than in Test 2 due to the compressed window leaving less time for the upper tank levels to equilibrate between transitions.

SC-MPC achieves the lowest total control energy (17 607 V^2 s) and the lowest total variation (84.8 V) among all four controllers, consistent with its energy-efficient behaviour across previous tests. The initial $0 \rightarrow 10$ cm settling is delayed to 36.3 s and 36.5 s on h_2 and h_4, the slowest among all controllers, as the scenario-based constraint tightening limits available control authority when both channels simultaneously demand large flow. The most notable observation in this test is the h_2 up-step overshoot of 31.84%, the highest single overshoot value recorded for SC-MPC across all tests. When both channels step up simultaneously, the combined flow demand pushes the optimiser against the tightened constraint margins on both outputs at the same time, temporarily reducing the degrees of freedom available to suppress the transient excursion on h_2. On h_4, the down-step overshoot reaches 29.76%, also elevated relative to Test 2. Steady-state accuracy remains competitive, with mean errors of 0.0895 and 0.0966 cm on h_2 and h_4.

The simultaneous excitation in Test 3 amplifies the cross-coupling effects relative to Test 2, and this is reflected in uniformly higher RMSE values on h_2 across all controllers. The ranking of controllers by IAE and energy remains consistent with previous tests: LQI leads on tracking speed, SL-MPC leads on h_4 accuracy and actuator smoothness, and SC-MPC leads on energy efficiency. The key difference from Test 2 is that SC-MPC's conservative constraint tightening, which was manageable under sequential excitation, becomes a more visible limitation when both channels compete for flow simultaneously, producing the largest overshoots of any controller in the incremental phases.

A comparison between Test 2 and Test 3 reveals that the simultaneous pulse structure does not significantly alter the initial $0 \rightarrow 10\,\text{cm}$ settling times but consistently increases overshoot on the incremental steps. This is because the upper tank levels h_1 and h_3 are driven harder when both pumps respond to simultaneous reference changes, creating a larger residual coupling disturbance at the moment the $10 \rightarrow 12$ and $12 \rightarrow 10\,\text{cm}$ transitions are applied. LQI and Nominal MPC, which do not explicitly model this coupling effect in their predictions, are more sensitive to this disturbance than SL-MPC, which partially compensates through its updated Jacobian.

5.5. Test 4: Differential (Asymmetric) Reference Tracking

Test 4 introduces an asymmetric reference profile in which the two controlled tanks receive opposing step changes after the initial filling phase. Specifically, h_2 follows $0 \rightarrow 10 \rightarrow 12 \rightarrow 10\,\text{cm}$ while h_4 follows $0 \rightarrow 10 \rightarrow 8 \rightarrow 10\,\text{cm}$. This differential excitation creates a demanding cross-coupling condition: one pump must increase flow while the other decreases it, placing the two channels in direct competition for actuator authority. Only the three MPC-based controllers are evaluated in this test, since LQI lacks the predictive decoupling capability required to handle conflicting reference commands without excessive cross-channel interference. The per-phase settling times and overshoot values in Table 5.9 complement the global metrics of Table 5.8 by isolating how each controller manages the opposing demands across the two channels.

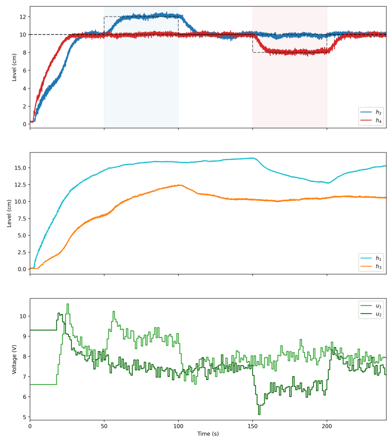

Figure 5.14: Nominal MPC

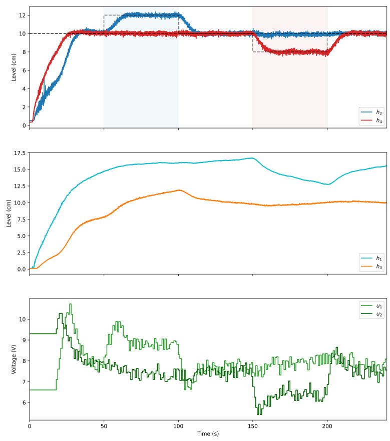

Figure 5.15: SL-MPC

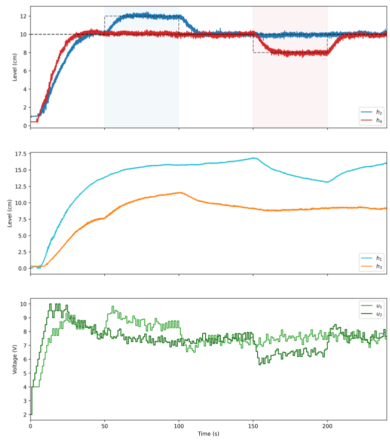

Figure 5.16: Closed-loop SC-MPC

Table 5.8: Global metrics for Test 4 (differential references).

	IAE (cm·s)		RMSE (cm)		SS Error (cm)		Energy / Variance	
Controller	h_2	h_4	h_2	h_4	h_2	h_4	Total	Var
SC-MPC	95.1	72.5	0.985	0.700	0.0969	0.3607	26 766	136.0
SL-MPC	98.2	54.4	1.092	0.528	0.0816	0.3609	27 473	126.4
Nominal MPC	102.2	56.2	1.079	0.530	0.0901	0.3414	27 636	127.8

Table 5.9: Transition metrics for Test 4 (differential references).

	0 → 10 cm (both)		h_2: 10 → 12; h_4: 10 → 8		h_2: 12 → 10; h_4: 8 → 10	
Controller	$T_s^{h_2}$ (s)	$T_s^{h_4}$ (s)	T_s / OS% (h_2)	T_s / OS% (h_4)	T_s / OS% (h_2)	T_s / OS% (h_4)
SC-MPC	36.8	29.8	10.7 / 23.59	11.8 / 27.97	12.2 / 20.57	11.1 / 13.59
SL-MPC	43.3	25.0	12.1 / 17.88	10.1 / 26.17	10.3 / 20.08	10.2 / 17.92
Nominal MPC	34.2	26.7	12.1 / 26.61	10.1 / 20.86	11.7 / 34.53	10.4 / 19.47

Table 5.10: Control effort comparison for Test 4 (differential references).

Controller	Total Energy	Input Variance	Voltage Range (V)
SC-MPC	26 766	136.0	2–10
SL-MPC	27 473	126.4	6–10.5
Nominal MPC	27 636	127.8	5–10

Nominal MPC achieves the fastest initial $0 \rightarrow 10\,\mathrm{cm}$ settling on h_2 (34.2 s) among the three controllers, and competitive settling on h_4 (26.7 s), consistent with its behaviour in previous tests where the fixed linear model performs adequately near the nominal operating point. On the differential phase, however, Nominal MPC exhibits the largest overshoot on the h_2 return step (34.53%), the highest single overshoot value recorded for any controller in this test. This reflects the heightened sensitivity to model mismatch when both channels operate simultaneously away from the nominal linearisation point in opposing directions: the fixed prediction model cannot accurately capture the coupled outflow dynamics under this asymmetric condition, causing the optimiser to misjudge the corrective action required on h_2 as it returns from 12 cm. Steady-state accuracy on h_2 is good (0.0901 cm mean error), but h_4 accumulates a persistent mean steady-state error of 0.3414 cm, significantly larger than in any symmetric test. Control energy ($27\,636\,\mathrm{V}^2\,\mathrm{s}$) and total variation (127.8 V) are the highest among the three controllers.

SL-MPC delivers the best h_4 performance in this test, achieving the lowest IAE (54.4 cm s), lowest RMSE (0.528 cm) and the fastest h_4 settling during the differential phase (10.1 s on both the up and down transitions). The online re-linearisation provides a measurable advantage here: as h_4 steps down to 8 cm while h_2 steps up to 12 cm, the updated Jacobian at each sample tracks the diverging operating points of the two channels more accurately than the fixed model, reducing the prediction error that drives cross-channel interference. On h_2, SL-MPC is slower during the initial fill (43.3 s), the slowest among the three controllers on this channel, but overshoot on the differential transitions is moderate (17.88% on the up-step and 20.08% on the return), lower than both Nominal MPC and SC-MPC on the h_2 return step. SL-MPC achieves the lowest total variation (126.4 V), preserving its characteristic actuator smoothness even under opposing reference demands. The persistent steady-state error on h_4 (0.3609 cm) is comparable to Nominal MPC, indicating that successive linearisation does not fully resolve the steady-state coupling disturbance introduced by the asymmetric flow routing.

SC-MPC achieves the lowest total control energy ($26\,766\,\mathrm{V}^2\,\mathrm{s}$) and the best global h_2 IAE (95.1 cm s) among the three controllers, consistent with its energy-efficient and h_2-competitive behaviour in previous tests. The initial $0 \rightarrow 10\,\mathrm{cm}$ settling is 36.8 s on h_2 and 29.8 s on h_4, with the h_4 settling being the fastest among all three controllers at this phase. On the differential transitions, SC-MPC shows elevated overshoot on h_4 during the down-step (27.97%), reflecting the difficulty of the scenario-based tightening when the two output constraints move in opposing directions: the robust margins on h_4 restrict the rate at which the controller can reduce that level while simultaneously allowing h_2 to rise. On the return phase, SC-MPC recovers well, with the lowest h_4 overshoot (13.59%)

of all three controllers. The persistent steady-state error on h_4 (0.3607 cm) is essentially identical to SL-MPC, confirming that this offset is a structural consequence of the plant topology rather than a controller-specific limitation.

A striking feature shared by all three controllers in Test 4 is the large persistent steady-state error on h_4 (approximately 0.34–0.36 cm across all controllers), which is an order of magnitude larger than the steady-state errors recorded in the symmetric tests. This behaviour is a direct consequence of the plant's cross-coupling structure: maintaining h_4 at 8 cm while h_2 is held at 12 cm requires Pump 2 to operate at a lower voltage than in the symmetric case, which simultaneously reduces the inflow to Tank 1 through the γ_2 split path. The resulting imbalance in the upper tank dynamics creates a residual coupling disturbance on h_4 that the velocity-form MPC formulation can partially but not fully reject within the experiment window. The fact that all three controllers converge to the same steady-state error level confirms that this offset is intrinsic to the plant's asymmetric flow routing and is not attributable to tuning differences between the controllers.

Comparing the three controllers overall, SL-MPC is the strongest option for h_4 tracking and actuator smoothness under differential excitation, while SC-MPC offers the best h_2 tracking and the lowest energy consumption. Nominal MPC provides competitive initial transient performance but is the most sensitive to model mismatch during the opposing differential transitions, as evidenced by the largest overshoot on the h_2 return step. The differential test therefore reinforces the conclusion that successive linearisation provides a tangible benefit whenever the two channels operate at significantly different levels simultaneously, and that the scenario-based approach remains effective for energy management even under the most demanding cross-coupling conditions.

5.6. Test 5: Maximum Limit Setpoint Tracking

Test 5 commands a single large step from empty to 15 cm on both controlled tanks simultaneously. This setpoint lies near the physical maximum of the apparatus and places the system in a nonlinear regime throughout the entire transient, since the Torricelli outflow gain varies substantially as the levels rise from zero to 15 cm. Because the tanks must be filled to a high level, the pumps operate at or near saturation for a sustained period, making this the most demanding test of large-signal behaviour, actuator-limit handling and model accuracy far from the nominal linearisation point. A critical experimental detail that affects the interpretation of all metrics is the difference in controller activation times: SC-MPC engages at $t = 1$ s, while SL-MPC and Nominal MPC activate at $t = 17$ s and LQI at $t = 19$ s. During the delay, the pumps run open-loop at their initial voltage, pre-filling the tanks to roughly 3–5 cm before the controller takes over. Controllers that activate earlier therefore accumulate tracking error over a longer controlled interval, which inflates their IAE and RMSE values relative to controllers that activate later and inherit a partially pre-filled system. Settling times in Table 5.12 are measured from controller activation to allow a fair transient comparison independent of this initialisation difference.

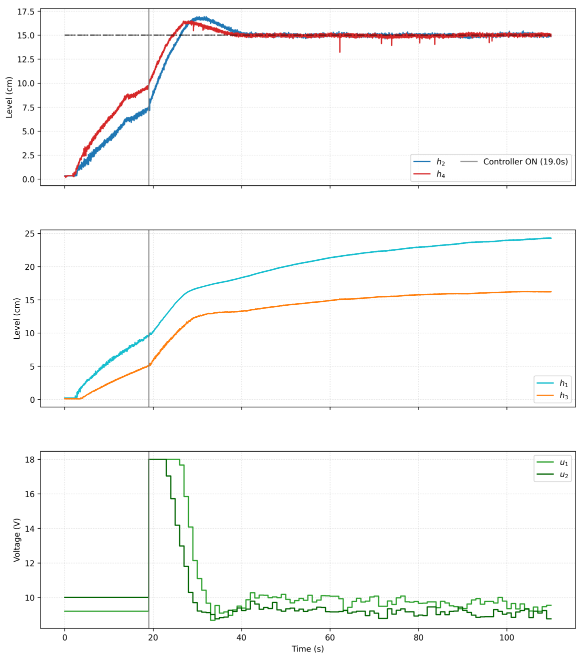

Figure 5.17: LQI

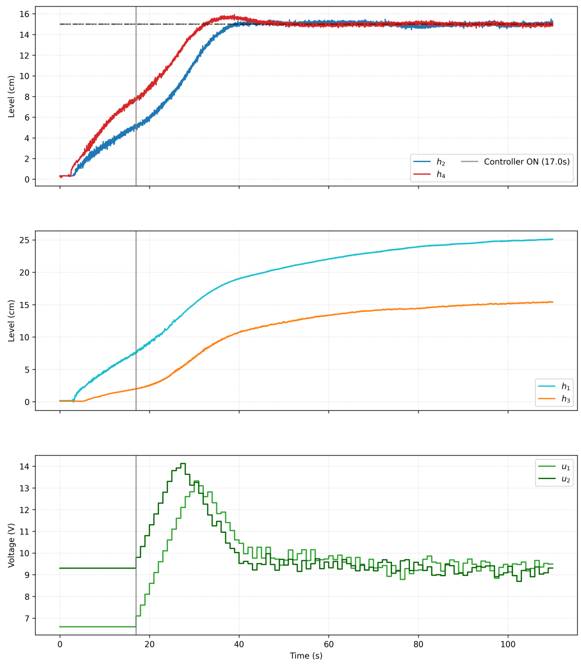

Figure 5.18: Nominal MPC

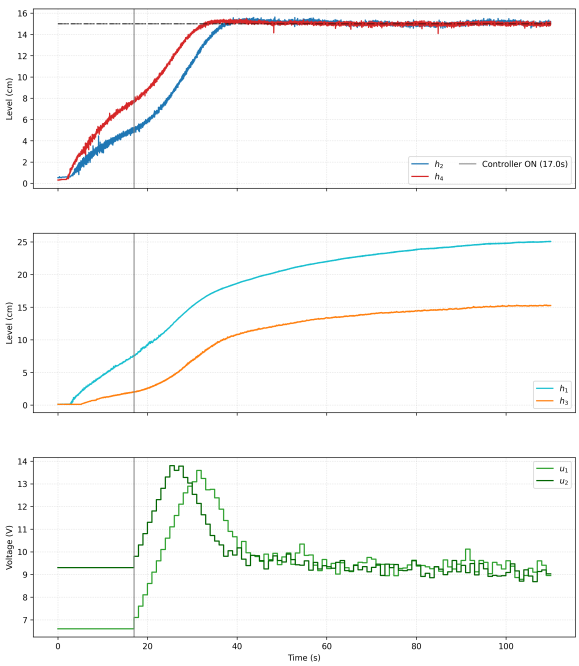

Figure 5.19: SL-MPC

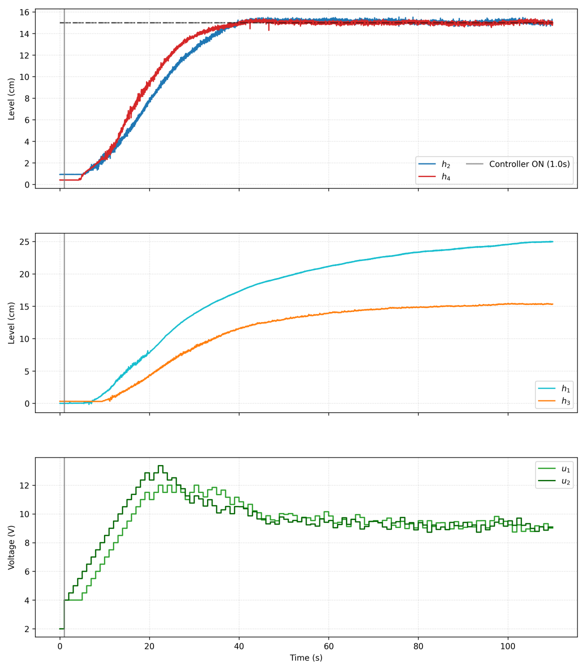

Figure 5.20: Closed-loop SC-MPC

Table 5.11: Global metrics for Test 5 ($0 \rightarrow 15$ cm maximum limit tracking).

	IAE (cm·s)		RMSE (cm)		SS Error (cm)	
Controller	h_2	h_4	h_2	h_4	h_2	h_4
LQI	255.7	208.9	4.856	4.201	0.0753	0.0864
SC-MPC	305.0	271.1	5.375	5.151	0.1274	0.0810
SL-MPC	324.2	245.7	5.516	4.608	0.0991	0.0724
Nominal MPC	326.7	253.0	5.574	4.686	0.1113	0.0753

Table 5.12: Time response for Test 5. Settling times are measured from controller activation.

	Ctrl ON	Settling Time (s)		Overshoot (%)	
Controller	(s)	h_2	h_4	h_2	h_4
LQI	19.0	37.6	34.3	13.22	10.12
SC-MPC	1.0	36.8	34.2	3.80	2.83
SL-MPC	17.0	35.5	30.5	4.38	3.40
Nominal MPC	17.0	36.5	39.8	3.36	6.32

Table 5.13: Control effort for Test 5.

Controller	Total Energy	Var (u_1)	Var (u_2)	Voltage Range (V)
LQI	23 525	33.7	29.7	8–18
SC-MPC	20 292	40.1	39.3	2–13
SL-MPC	20 636	30.1	27.2	7–14
Nominal MPC	20 895	30.9	27.3	7–14

LQI records the lowest IAE (255.7 and 208.9 cm s on h_2 and h_4) and the lowest RMSE (4.856 and 4.201 cm) among all four controllers, but this advantage must be interpreted in the context of its late activation at $t = 19$ s: the pre-filling phase brings the tanks to approximately 3–5 cm before the controller engages, effectively reducing the error accumulated during the most nonlinear part of the fill. From activation, LQI settles in 37.6 s on h_2 and 34.3 s on h_4, competitive but not the fastest in this test. The primary drawback is the aggressive actuation required to drive the system to 15 cm: pump voltage saturates at 18 V immediately after controller engagement and remains elevated throughout the transient, producing the highest overshoot of all controllers (13.22% on h_2 and 10.12% on h_4) and the highest total energy (23 525 V^2 s). The lack of explicit constraint handling means LQI cannot moderate its actuation as the output approaches the setpoint, resulting in the largest transient excursion observed in any test.

Nominal MPC activates at $t = 17$ s and achieves the lowest overshoot on h_2 (3.36%), demonstrating that explicit constraint handling effectively suppresses aggressive actuation even near the physical maximum. However, this test exposes the clearest degradation associated with the fixed linear prediction model: the settling time on h_4 increases to 39.8 s, the slowest of any controller on any channel in this test, while the settling time on h_2 is 36.5 s. As the tanks fill toward 15 cm, the linearisation error grows with the distance from the nominal operating point at 10 cm, causing the prediction model to underestimate the outflow resistance and produce suboptimal control sequences during the upper portion of the fill transient. The h_4 overshoot of 6.32% is also the highest among the three MPC controllers, further reflecting the accumulated prediction error at high levels. Control energy (20 895 V^2 s) and per-pump variation (30.9 and 27.3 V) are mid-range, with voltage remaining within 7–14 V throughout.

SL-MPC activates at $t = 17$ s and delivers the best overall transient performance among all four controllers, achieving the fastest settling times from activation on both channels (35.5 s on h_2 and 30.5 s on h_4) while keeping overshoot low (4.38% and 3.40%). The advantage over Nominal MPC is directly attributable to the online re-linearisation: as the tank levels rise continuously from 3–5 cm to 15 cm, the updated Jacobian at each sample tracks the changing Torricelli gain and maintains prediction accuracy throughout the wide operating range traversed during this test. This is precisely the scenario where successive linearisation provides its greatest benefit, as the linearisation error of a fixed model grows rapidly with the distance from the 10 cm design point. SL-MPC achieves the lowest per-pump variation (30.1 and 27.2 V) and operates entirely within 7–14 V, confirming smooth actuation without saturation. Total energy (20 636 V^2 s) is slightly higher than SC-MPC but comparable. Steady-state accuracy is strong, with mean errors of 0.0991 and 0.0724 cm on h_2 and h_4.

SC-MPC activates earliest at $t = 1$ s, which directly explains its highest IAE values (305.0 and 271.1 cm s) and highest RMSE (5.375 and 5.151 cm): the controller accumulates error across the full nonlinear fill from near-zero levels, where the linearisation error is largest and the scenario-based constraint tightening is most restrictive. Despite this, SC-MPC achieves the lowest total energy (20 292 V^2 s) and the most effective overshoot suppression of all controllers (3.80% on h_2 and 2.83% on h_4), demonstrating that the robust constraint tightening actively limits aggressive actuation even under the most demanding conditions.

The settling times from activation (36.8 s on h_2 and 34.2 s on h_4) are competitive with LQI despite starting from near-empty tanks, which is notable given the conservative constraint margins. The per-pump variation (40.1 and 39.3 V) is the highest among the MPC controllers, driven by more frequent corrective actions as the scenario bounds constrain the available input authority near the tightened output limits. Voltage range extends from 2 to 13 V, the widest of any controller in this test.

Test 5 provides the clearest illustration of the trade-off between model accuracy and constraint handling across the full controller hierarchy. LQI achieves competitive settling and the best IAE metrics, but only because it activates late and inherits a partially pre-filled system, and its aggressive actuation produces the largest overshoot and highest energy. Among the MPC controllers, SL-MPC offers the best compromise: it benefits from the pre-filling phase in the same way as Nominal MPC but adds the accuracy of online re-linearisation to achieve faster settling and lower overshoot on both channels. Nominal MPC shows the clearest performance degradation of any controller in this test, with the slowest h_4 settling and the highest h_4 overshoot among the MPC variants, confirming that the fixed linearisation model is most penalised when the operating range is widest.

The consistently low overshoot values of all three MPC controllers (2.83–4.38% on both channels) compared to LQI (10.12–13.22%) demonstrate the practical value of explicit constraint handling when operating near the physical limits of the apparatus. Without the MPC optimisation enforcing the level bounds over the prediction horizon, LQI has no mechanism to anticipate the approach to the 15 cm setpoint and moderate its actuation accordingly. This distinction becomes particularly important in safety-critical applications where overshoot near a physical boundary could trigger overflow conditions.

5.7. Asymmetric Setpoint Feasibility Analysis

The multi-step reference tests in the preceding sections use setpoints within a safe operating window. This section documents the practical feasibility boundaries of the quadruple tank system, which constrain the admissible reference combinations for all controllers.

The feasible setpoint range for each controlled output was determined empirically by computing the steady-state equilibrium $(\bar{x}, \bar{u})$ for a grid of reference values and checking whether the solution satisfies the hardware constraints: pump voltages within $[4, 18]$ V and all four tank levels within $[0, 27]$ cm. The resulting operating window is

$$y_{\min} = 9 \text{ cm} \; \leq \; h_{2,\text{ref}}, \; h_{4,\text{ref}} \; \leq \; 15 \text{ cm} = y_{\max}. \tag{5.7}$$

Below 9 cm the pumps approach their minimum effective voltage threshold (dead-zone region of the pump map), making the operating point unreliable. Above 15 cm the uncontrolled tanks approach the physical overflow limit, as the equilibrium levels in the upper tanks grow rapidly with the controlled output references.

The cross-coupling structure of the quadruple tank imposes a directional constraint on asymmetric setpoints. From the mass balances (3.5), maintaining a high level in Tank 4 requires large flow through the $(1-\gamma_2)$ path of Pump 2. However, a fraction γ_2 of the same pump flow is simultaneously directed to Tank 1. If $h_{4,\text{ref}}$ is high while $h_{2,\text{ref}}$ is low, Tank 1 receives substantial inflow from Pump 2 but drains slowly (since Tank 2 is kept at a low level, Pump 1 operates at moderate voltage and the $h_1 \to h_2$ drainage path is the only outlet for Tank 1). This imbalance causes the equilibrium level $\bar{h}_1$ to exceed the physical tank height.

Concretely, for the configuration $h_{2,\text{ref}} = 9\,\text{cm}$, $h_{4,\text{ref}} = 12\,\text{cm}$ (denoted Configuration A), the equilibrium solver returns $\bar{h}_1 \approx 30\,\text{cm}$, which exceeds the 27 cm safe limit. The experimental test confirms this: Tank 1 accumulates water continuously until the safety limit triggers experiment termination.

The reversed configuration $h_{2,\text{ref}} = 12\,\text{cm}$, $h_{4,\text{ref}} = 9\,\text{cm}$ (Configuration B) is feasible. Both LQI and Nominal MPC track the asymmetric references without instability, and the uncontrolled tanks remain bounded.

This analysis establishes the constraint

$$h_{2,\text{ref}} \;\geq\; h_{4,\text{ref}} \tag{5.8}$$

as a necessary condition for feasible steady-state operation. The constraint arises from the physical flow routing and is independent of the controller architecture.

Within the feasible region defined by (5.7) and (5.8), the following asymmetric setpoint combinations were validated through equilibrium computation and closed-loop simulation. All combinations produce bounded equilibria for all four tanks and pump voltages within the hardware limits.

Table 5.14: Validated asymmetric setpoint combinations. All satisfy $h_{2,\text{ref}} \geq h_{4,\text{ref}}$ and produce feasible equilibria within hardware constraints.

Setpoint	$h_{2,\text{ref}}$ (cm)	$h_{4,\text{ref}}$ (cm)
SP1	10	9
SP2	11	9
SP3	12	10
SP4	13	10
SP5	14	11

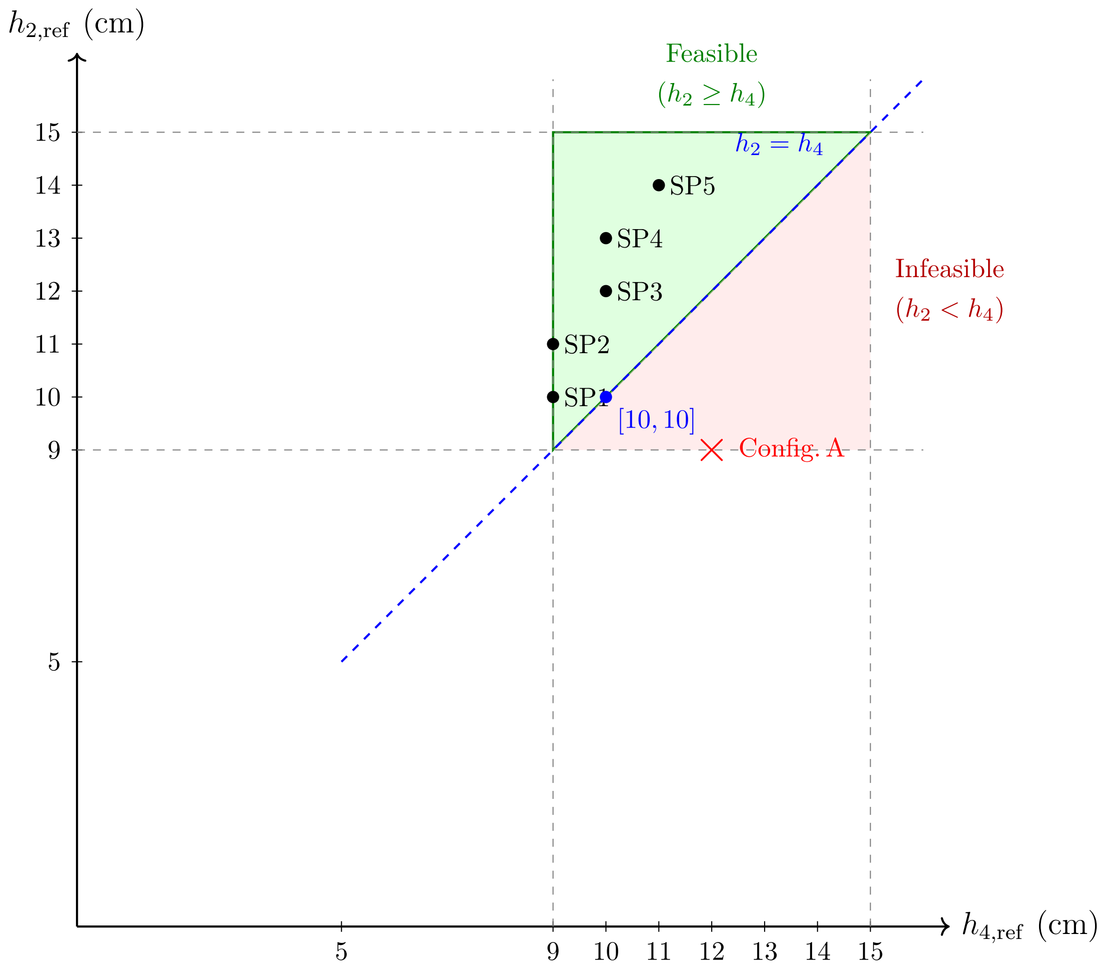

Figure 5.21: Feasibility map of the admissible reference space. The green region satisfies both the operating range constraints $9 \leq h_{2,\text{ref}}, h_{4,\text{ref}} \leq 15\,\text{cm}$ and the asymmetric condition $h_{2,\text{ref}} \geq h_{4,\text{ref}}$. Validated setpoints are marked with filled circles; the infeasible Configuration A is marked with a cross.

Table 5.15 collects the empirically determined operating limits.

Table 5.15: Operating region constraints for the quadruple tank system.

Constraint	Value	Origin
Minimum feasible setpoint	$\approx 9\,\mathrm{cm}$	Pumps approach dead-zone threshold
Maximum feasible setpoint	$\approx 15\,\mathrm{cm}$	Upper tanks approach overflow
Asymmetric feasibility	$h_{2,\mathrm{ref}} \geq h_{4,\mathrm{ref}}$	Tank 1 mass balance (cross-flow)
Tank overflow limit	$27\,\mathrm{cm}$	Physical tank height

These constraints are intrinsic to the plant configuration and flow routing, not to the choice of controller. Any supervisory layer or reference governor should enforce the feasible region (5.7)–(5.8) to prevent unsafe operating points.

5.8. Disturbance Rejection: Manual Water Injection

This test evaluates the controllers' ability to reject unmeasured disturbances while maintaining a 10 cm setpoint. Once steady-state is reached, manual water injections are applied directly into the controlled tanks at two instants: at $t \approx 100\,\mathrm{s}$ into Tank 2, with analysis window $[100, 140]\,\mathrm{s}$, and at $t \approx 150\,\mathrm{s}$ into Tank 4, with analysis window $[150, 190]\,\mathrm{s}$. Because the injections are performed by hand, the disturbance magnitude is not perfectly repeatable across controllers. The segmented metrics in Table 5.17 therefore provide the most meaningful basis for comparison, since they isolate the local rejection behaviour within each injection window independently of the global tracking performance reported in Table 5.16. The global metrics are dominated by the initial $0 \to 10\,\mathrm{cm}$ fill transient and are included for completeness, but the segmented peak deviation and recovery time values are the primary indicators of disturbance rejection quality in this test.

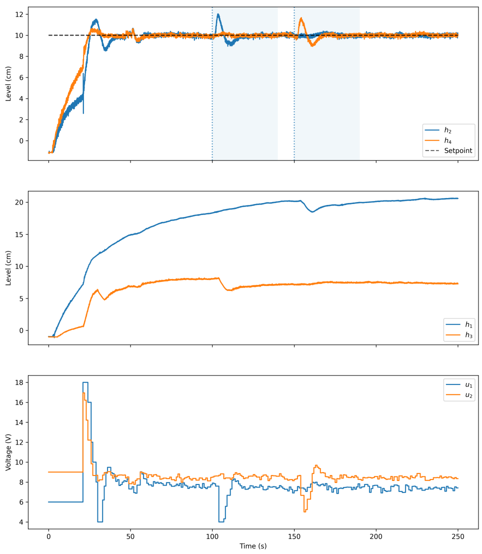

Figure 5.22: LQI with 15s open-loop startup

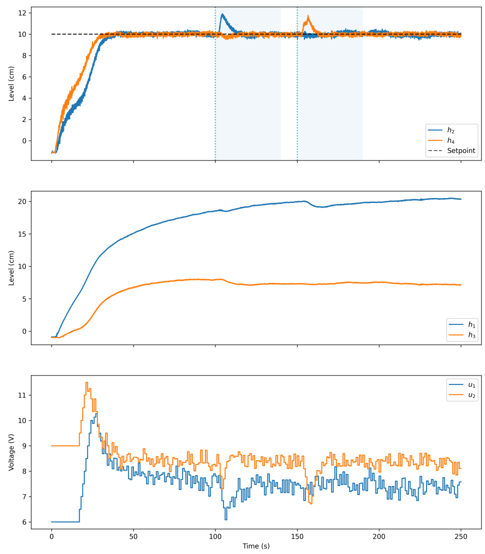

Figure 5.23: Nominal MPC with 15s open-loop startup

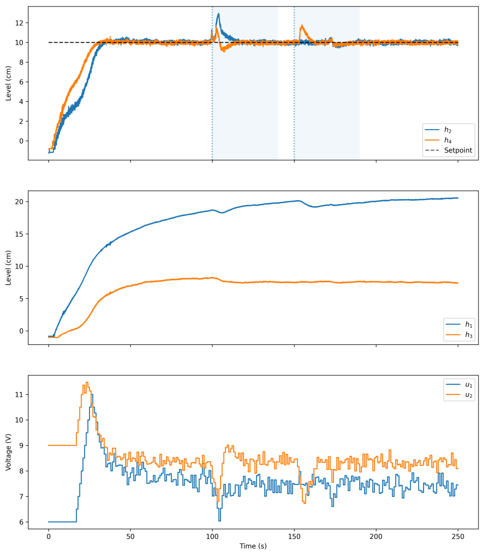

Figure 5.24: SL-MPC with 15s open-loop startup

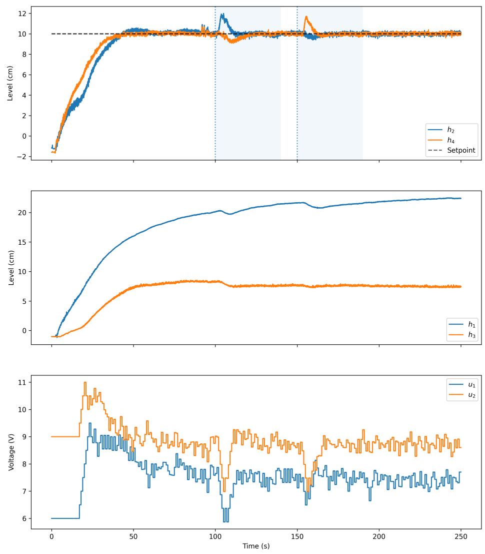

Figure 5.25: SC-MPC with 15s open-loop startup

Table 5.16: Global metrics for the disturbance rejection test (10 cm setpoint with manual injections).

	IAE (cm·s)		RMSE (cm)		SS Error (cm)		Energy / Variance	
Controller	h_2	h_4	h_2	h_4	h_2	h_4	Total	Var
LQI	90.7	55.4	1.141	0.676	0.0683	0.0697	31 960	133.5
SC-MPC	114.1	78.1	1.287	0.882	0.0719	0.0613	32 151	140.7
SL-MPC	96.7	59.8	1.199	0.676	0.0785	0.0677	30 745	122.0
Nominal MPC	94.4	56.7	1.179	0.707	0.0937	0.0696	30 700	134.6

Table 5.17: Segmented disturbance response metrics. Peak = maximum deviation from setpoint. T_r = recovery time.

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

Table 5.18: Control effort for the disturbance rejection test.

Controller	Total Energy	Input Variance	Voltage Range (V)
LQI	31 960	133.5	4–18
SC-MPC	32 151	140.7	6–11
SL-MPC	30 745	122.0	6–11
Nominal MPC	30 700	134.6	6–11

LQI achieves the lowest global IAE and RMSE on both channels (90.7 and 55.4 cm s, 1.141 and 0.676 cm), confirming the strong nominal tracking performance observed across all previous tests. Within the injection windows, LQI records the largest peak deviations of any controller on both disturbances: 2.031 cm on the h_2 injection and 1.674 cm on the h_4 injection. Recovery times are also the longest at 15.8 s and 15.4 s respectively. This behaviour is consistent with the LQI architecture: without a predictive model to anticipate the propagation of the disturbance through the coupled tank dynamics, the integral action can only respond reactively once the output error has accumulated, resulting in a larger initial excursion and a slower return to steady state. The total energy of 31 960 V^2 s and total variation of 133.5 V are the highest among all controllers, with pump voltage reaching 18 V during the disturbance rejection transients, confirming the aggressive actuation characteristic of LQI under perturbation.

Nominal MPC achieves competitive global metrics (94.4 and 56.7 cm s IAE, 1.179 and 0.707 cm RMSE) and delivers balanced disturbance rejection performance in both injection windows. The peak deviation on the h_2 injection is the lowest of all controllers at 1.905 cm, demonstrating that the MPC cost function effectively moderates the initial excursion by penalising aggressive actuation in the receding horizon. Recovery time is 11.8 s on h_2 and 11.5 s on h_4, moderate compared to SC-MPC but substantially faster than LQI. The steady-state error remains low at 0.0937 cm on h_2 and 0.0696 cm on h_4, confirming that the velocity-form formulation rejects the constant-level disturbance without persistent offset. Control energy (30 700 V^2 s) and total variation (134.6 V) are mid-range, with voltage remaining within 6–11 V throughout the experiment.

SL-MPC records competitive global metrics (96.7 and 59.8 cm s IAE) but shows the largest peak deviation on the h_2 injection (2.957 cm), the highest of any controller in this test. This result is notable given SL-MPC's strong performance in the tracking tests and warrants explanation. When the manual injection abruptly raises h_2 away from the steady-state operating point, the successive linearisation updates the Jacobian at the new perturbed level, but the prediction model requires one or two samples to stabilise around the correct linearisation after the sudden state change. During this brief adjustment, the control action is computed from a prediction model that is catching up to the true operating point, producing a larger transient excursion than controllers that use a fixed model and are therefore not subject to this re-linearisation lag. Recovery time on h_2 is 13.2 s, faster than LQI but slower than SC-MPC and Nominal MPC. On the h_4 injection, the re-linearisation lag is less pronounced since the h_4 disturbance is smaller and the Jacobian update converges more quickly, resulting in a peak deviation of 1.805 cm and a recovery time of 10.2 s, both competitive. SL-MPC achieves the lowest total variation (122.0 V) and the lowest total energy (30 745 V^2 s), confirming its characteristic actuator smoothness is maintained even under disturbance conditions.

SC-MPC recovers fastest from both disturbances, with the shortest recovery times of any controller: 10.3 s on the h_2 injection and 10.7 s on the h_4 injection. The segmented IAE values are also the lowest in both windows (11.75 and 8.62 cm s), demonstrating that the scenario-based robust optimisation is particularly effective at localising the disturbance response and returning quickly to the reference. This advantage arises because the scenario-tightened constraints maintain feasibility margins around the current operating point even under perturbation, giving the optimiser sufficient room to apply a strong corrective action without violating the output bounds. The peak deviations (2.000 cm on h_2 and 1.715 cm on h_4) are moderate, between LQI and Nominal MPC on h_2 and comparable to Nominal MPC on h_4. The global IAE on h_2 (114.1 cm s) and RMSE (1.287 cm) are the highest among all controllers, driven by the conservative constraint tightening during the initial fill rather than by poor disturbance rejection. Total energy (32 151 V^2 s) and total variation (140.7 V) are the highest among the MPC controllers in this test, consistent with the more frequent corrective actions required by the robust constraint handling near the tightened bounds.

The disturbance rejection test reveals a ranking inversion relative to the tracking tests. LQI, which dominated on global IAE across most tracking scenarios, performs worst on the segmented rejection metrics, recording the largest peak deviations and the slowest recovery times in both injection windows. SC-MPC, which was the most conservative controller during large transients, demonstrates the strongest local rejection capability, exploiting its robust optimisation structure to return to steady state faster than any other controller. This inversion reflects a fundamental architectural difference: LQI responds to disturbances only after the error has accumulated in the integral state, while SC-MPC anticipates the disturbance propagation through the coupled tank dynamics within the prediction horizon and applies corrective action earlier in the transient.

A secondary observation is that the h_2 injection consistently produces larger peak deviations and longer recovery times than the h_4 injection across all controllers. This asymmetry is explained by the plant's cross-coupling structure: a disturbance injected into Tank 2 propagates through both the direct h_2 control loop and the indirect path via Tank 1, while a disturbance into Tank 4 propagates primarily through the direct h_4 loop with weaker coupling to the rest of the network. The consistently lower peak deviations and faster recovery times on h_4 across all controllers confirm that this asymmetry is a plant property rather than a controller-dependent effect.

5.9. Comparison of SC-MPC and Adaptive SC-MPC

To demonstrate the practical benefit of online learning, both the SC-MPC (Section 4.3) and the adaptive SC-MPC are tested on the real quadruple tank plant under identical conditions using the setpoint tracking experiment.

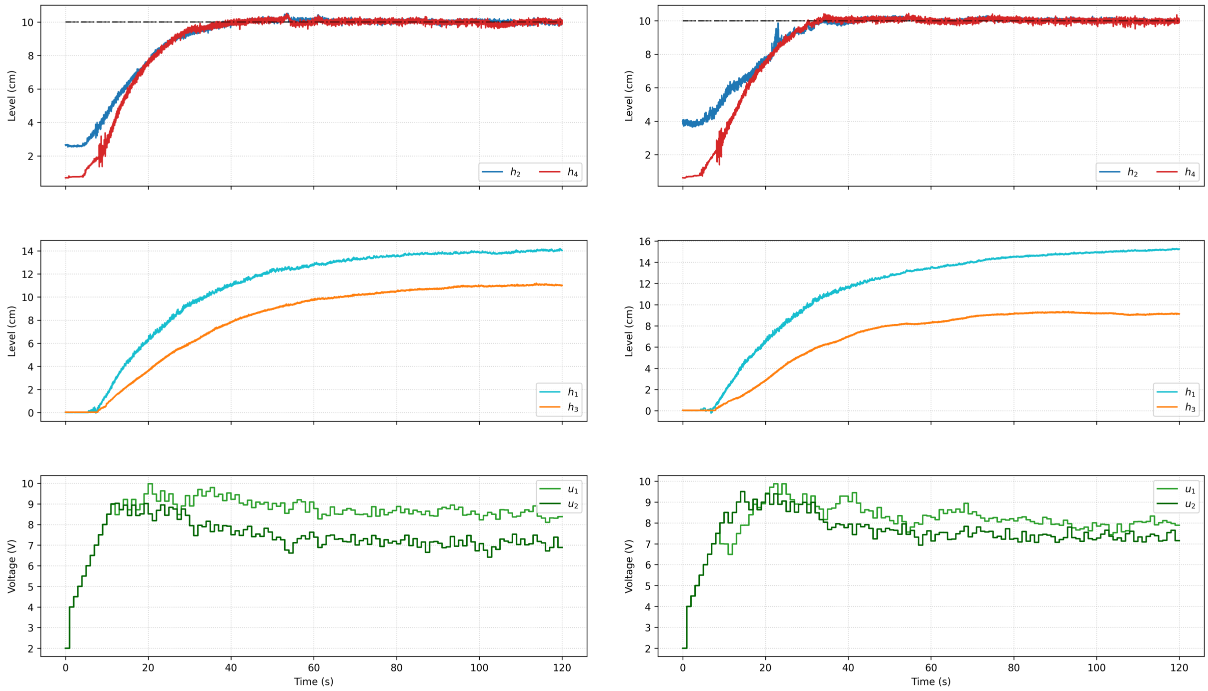

(a) SC-MPC results. (b) Adaptive SC-MPC results.

Figure 5.26: Setpoint tracking on the real plant.

Table 5.19 compares the key performance metrics from both experiments. The adaptive controller achieves lower integrated absolute error and RMSE on both outputs, reduced overshoot, and lower control energy and total variation compared to the fixed SC-MPC.

Table 5.19: Performance comparison: SC-MPC vs adaptive SC-MPC (Test 1, real plant).

Metric	**SC-MPC**	**Adaptive SC-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 time h_2 (s)	37.1	38.4	—
Settling time h_4 (s)	35.4	30.0	15%
Control energy (V^2·s)	15523	14955	4%
Total variation (V)	73.6	67.8	8%

The improvement is attributed to the contraction of the uncertainty box through online learning. During the first 40 samples the SMI is inactive and both controllers behave identically (wide box, conservative tightening). After the guard period expires, the adaptive controller begins learning the valve coefficients from steady-state measurements. As the box shrinks, the scenario spread narrows, the constraint margins decrease, and the QP optimizer gains additional freedom to minimize the tracking cost. This manifests as reduced overshoot (the controller can apply more aggressive corrections without violating the tightened bounds) and lower control energy (less conservative actuation is needed to maintain constraint satisfaction).

The settling time on h_2 is marginally slower for the adaptive controller (38.4 s vs 37.1 s), which falls within the experimental repeatability margin and is not considered significant. On h_4, the adaptive controller settles 5.4 s faster, reflecting the benefit of tighter uncertainty quantification on the cross-coupled output.

The adaptive SC-MPC represents the most capable controller in the comparison study. It combines the constraint handling of MPC, the robustness of the scenario approach, the accuracy of successive linearization and the learning capability of set membership identification. As the parameter box shrinks through online learning, the controller progressively transitions from conservative robust behaviour to near-nominal performance.

Several design choices affect the practical performance. The noise bound ε_{smi} should be set above the actual sensor noise level; too small a value risks excluding the true parameter after a noise spike, while too large a value slows learning. The expansion factor λ_{exp} provides a forgetting mechanism that enables the controller to track slowly drifting parameters. However, the frequency of expansion is as important as its magnitude: applying expansion at every timestep, even with a small factor, actively counteracts the contraction from strip intersections and prevents the box from shrinking below the initial range. The implementation applies expansion only once every 60 steps with $\lambda_{\text{exp}} = 0.005$, which provides sufficient collapse prevention while allowing genuine uncertainty reduction to accumulate between expansion events.

First, note that the maximum scenario envelope width provides a direct measure of SMI convergence at each timestep, shown in Figure 5.27.

The maximum scenario envelope width over the $N = 60$ step horizon contracts from an initial peak of approximately 0.9 cm to a steady-state of approximately 0.25 cm within the first 60 seconds, representing a 70% reduction driven by the first valid SMI observations after the guard period expires. Under continuous steady-state operation the scenario spread continues to contract slowly, with a further 10–11% reduction between the first and second halves of the 2000 s experiment. When pulse reference changes are applied the cooldown mechanism suspends learning during each transition, reducing the effective

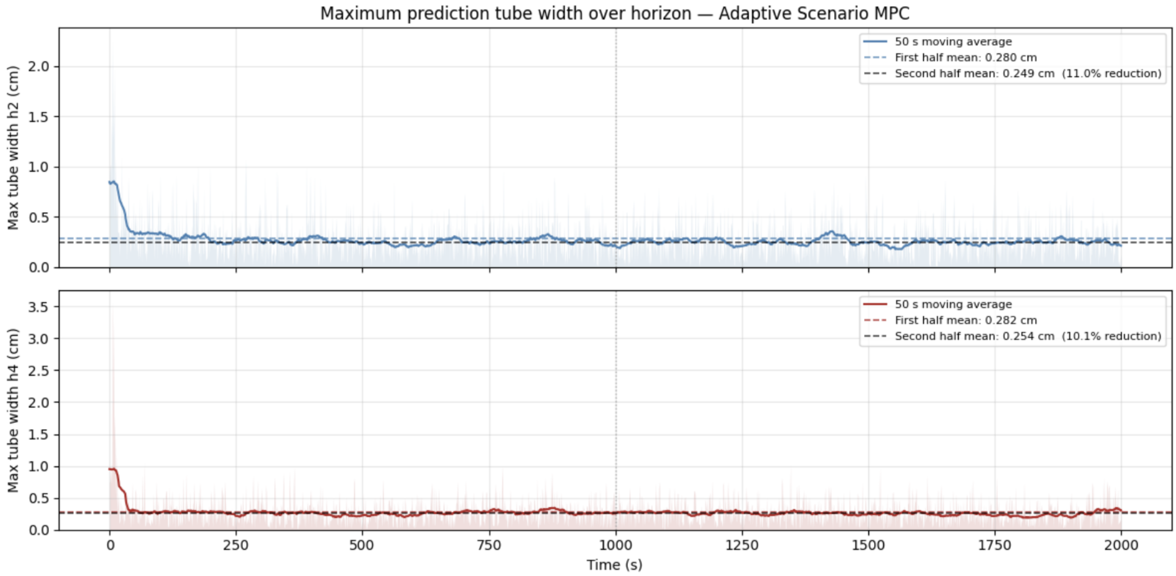

Figure 5.27: Maximum scenario envelope width over the $N = 60$ step prediction horizon for h_2 (top) and h_4 (bottom) during a 2000 s simulation.

number of updates and limiting the residual slow contraction to approximately 5–9%. The absolute steady-state scenario spread of 0.25 cm is small relative to the operating range and constraint margins, confirming that constraint tightening remains effective without being overly conservative throughout the experiment.

The reference-change cooldown is the most critical safety mechanism. During reference transitions, the levels change rapidly and the finite-difference derivative in (4.39) produces biased estimates of A_i^{obs}. Without the cooldown, the SMI interprets these transient dynamics as valve coefficient changes and contracts the box around incorrect values, leading to degraded or unstable closed-loop performance. The 100-sample cooldown ensures that learning resumes only after the system has settled at the new operating point, where the mass balance inversion is reliable. This approach was validated experimentally: disabling the cooldown caused controller instability during pulse tracking tests, while enabling it restored stable operation with progressive learning between setpoint changes.

The choice to anchor the nominal prediction model to the identified valve values A^{id} rather than the learned box center θ_{center} ensures that the prediction quality is decoupled from the SMI state. The learned box affects only the scenario generation and constraint tightening, providing a clean separation between the prediction model (which should reflect the best available plant knowledge) and the uncertainty quantification (which adapts online).

The scenario envelope evolution in Figure 5.27 was obtained in simulation using the nonlinear grey-box model of the quadruple-tank system with additive measurement noise

($\varepsilon_{\text{smi}} = 0.35\,\text{cm}$), reproducing the closed-loop behaviour observed on the physical plant. The comparison results in Figure 5.26 and Table 5.19, by contrast, were obtained directly on the laboratory plant.

The safety mechanisms described in Section 4.4 were not part of the initial adaptive SC-MPC design but were introduced iteratively in response to failures observed during preliminary laboratory experiments. In early tests without the reference-change cooldown or the derivative threshold, the controller operated correctly at steady state but lost feasibility shortly after each setpoint change. The root cause was traced to the mass balance inversion (4.39): during the transient following a reference step, the rapidly varying levels produced large finite-difference derivatives that biased the observed valve coefficients A_i^{obs}. The SMI interpreted these biased observations as valid and contracted the parameter box around incorrect values. Because the scenario-based constraint tightening is computed from the box, the tightened output bounds shifted away from the true feasible region, and the QP became infeasible within a few samples of the reference change. Introducing the 100-sample cooldown (n_{cool}) to inhibit learning during transients, complemented by the per-tank derivative threshold ($|\dot{h}_i| < 0.3\,\text{cm/s}$) as an additional guard against residual dynamics, eliminated the infeasibility events entirely. All results reported in Figure 5.26 and Table 5.19 were obtained with the full set of safety mechanisms active.

5.10. Summary of Findings

This section synthesises the experimental findings across all seven test scenarios into a unified comparative assessment of the five control strategies. The discussion is organised around four dimensions that are relevant to practical controller selection: tracking performance, constraint handling and actuator efficiency, robustness to nonlinearity and parametric uncertainty, and disturbance rejection capability.

Tracking Performance Across all symmetric tracking tests, LQI consistently achieves the lowest global IAE and the fastest settling times. This advantage is most pronounced in Test 1, where LQI settles h_2 in 25.8 s compared to 35.5 s for Nominal MPC and 49.0 s for SC-MPC. The same pattern repeats in Tests 2 and 3, where LQI leads on combined IAE in both channels despite the increased cross-coupling load. The speed advantage of LQI is structurally rooted in its unconstrained integral action, which applies the full available actuator authority immediately after each reference change without the conservative margins introduced by predictive constraint handling.

Among the MPC controllers, SL-MPC consistently delivers the best h_4 tracking across all multi-step tests. In Test 2, SL-MPC achieves the lowest h_4 IAE (55.8 cm s) and the fastest h_4 settling on the incremental steps (8.9 s), advantages that persist through Tests 3 and 4. This channel-specific benefit is directly attributable to the online re-linearisation, which tracks the level-dependent Torricelli gain as the operating point moves and maintains prediction accuracy in a way that the fixed nominal model cannot match. On h_2, however, SL-MPC consistently records higher IAE than Nominal MPC and SC-MPC during simultaneous filling from empty, reflecting the coupling asymmetry between the two channels when both are driven from zero simultaneously.

Nominal MPC provides the most balanced cross-channel tracking behaviour among the MPC controllers, with comparable IAE values on h_2 and h_4 across all tests. Its primary limitation is a systematic slowdown during large transients, particularly on h_2 in Tests 2 and 3, where the fixed linear model loses prediction accuracy as the operating point moves far from the 10 cm nominal linearisation. This degradation is most visible in Test 5, where the settling time on h_4 increases to 39.8 s, the slowest of any controller on any channel in the maximum-limit test.

SC-MPC consistently shows the most conservative transient behaviour, with settling times on h_2 ranging from 49.0 s in Test 1 to 36.8 s in Test 4, always at the slower end of the distribution. This conservatism is the direct consequence of the scenario-based constraint tightening, which reserves output margins for all $K = 39$ parameter realisations simultaneously and limits the control authority available during large transients. In the differential test (Test 4), SC-MPC recovers to offer the best h_2 IAE among the MPC controllers, suggesting that the robust tightening provides a structural advantage when cross-coupling creates asymmetric demands on the two channels.

Constraint Handling and Actuator Efficiency LQI handles constraints through post-hoc saturation only, with no mechanism to anticipate future constraint activity. As a result, pump voltage saturates at 18 V in every test involving a large initial step, and overshoot is consistently the highest of all controllers. In Test 5, LQI produces 13.22% overshoot on h_2, more than three times the overshoot of any MPC controller on the same channel. The total energy and total variation are also consistently the highest among all configurations with open-loop startup, confirming that the speed advantage of LQI is purchased through greater actuator stress.

All three MPC controllers enforce the level and voltage bounds through the QP optimisation and produce significantly lower overshoot in every test. SL-MPC achieves the lowest total variation across all tests, with values of 56.2 V in Test 1, 125.5 V in Test 2 and 76.1 V in Test 3, consistently below both Nominal MPC and SC-MPC. This smoothness benefit arises from the slew-rate penalty on Δu combined with the updated prediction model, which reduces the need for abrupt corrective actions by anticipating the operating-point-dependent dynamics more accurately.

SC-MPC achieves the lowest total energy in every test in which it participates, ranging from $14\,954\,\mathrm{V}^2\,\mathrm{s}$ in Test 1 to $20\,292\,\mathrm{V}^2\,\mathrm{s}$ in Test 5. The scenario-based optimisation finds control sequences that are globally less energetic by enforcing constraint satisfaction across all uncertainty realisations simultaneously, which implicitly prevents the large corrective spikes that drive energy consumption in the other controllers. This energy advantage comes at the cost of higher total variation than SL-MPC in several tests, reflecting the more frequent small corrective actions required near the tightened output bounds.

The RMSE safety margin of 2.09 cm applied to the effective level bounds confirms its practical necessity: across all tests involving large transients, none of the effective level bounds $[2.19, 24.91]$ cm were violated by either h_2 or h_4 under SC-MPC or adaptive SC-MPC. This empirical absence of violations, achieved under conditions specifically designed to stress the linear prediction model (large transients, near-saturation inputs, coupling-induced cross-disturbances), provides strong evidence that the chosen safety margin is both necessary and sufficient for the operating conditions considered.

Robustness to Nonlinearity and Parametric Uncertainty The tests involving wide operating ranges (Test 5) and opposing reference demands (Test 4) provide the clearest evidence of how each controller responds to nonlinearity and model mismatch. Nominal MPC is the most sensitive: its settling time degrades most significantly in Test 5 (39.8 s on h_4 compared to 25.8 s in Test 1) and it produces the largest overshoot on the h_2 return step in Test 4 (34.53%), both consequences of the growing linearisation error as the operating point moves away from the 10 cm nominal design.

SL-MPC mitigates this sensitivity effectively whenever a single channel traverses a wide operating range, as demonstrated by its consistent h_4 advantage in Tests 2, 3 and 4 and its fastest settling in Test 5 (35.5 s on h_2 and 30.5 s on h_4 from activation). The benefit is less pronounced when both channels fill simultaneously from empty, where the coupled dynamics introduce a re-linearisation lag that temporarily increases tracking error on h_2.

SC-MPC provides structured robustness to the valve coefficient uncertainty through the scenario-based tightening, which ensures constraint satisfaction with at least 95% probability under the assumed parametric uncertainty range. The fixed uncertainty set $[0.75, 0.90]\,A^{\mathrm{id}}$ is conservative by design, reflecting the offline worst-case characterisation of valve drift. As a consequence, SC-MPC maintains safe constraint margins even when the true valve coefficients differ from their nominal values, but at the cost of the consistent settling delays observed across all tests.

The adaptive SC-MPC addresses this conservatism through online set membership identification of the valve coefficients. As reported in Table **??**, the adaptive variant reduces IAE by 15% on h_2 and 3% on h_4, overshoot by 26% on h_2 and 15% on h_4, and total energy by 4% relative to the fixed SC-MPC under identical experimental conditions. These improvements are achieved as the parameter box contracts through online learning, progressively narrowing the scenario spread and reducing the constraint tightening margins. The scenario envelope width contracts from an initial peak of approximately 0.9 cm to a steady-state of 0.25 cm within the first 60 s, a 70% reduction that directly translates into less conservative output bounds and greater optimiser freedom to minimise the tracking cost.

Disturbance Rejection The disturbance rejection test reveals a ranking inversion relative to the tracking tests. LQI, which dominated on global IAE across most tracking scenarios, performs worst on the segmented rejection metrics, recording the largest peak deviations (2.031 cm on h_2 and 1.674 cm on h_4) and the slowest recovery times (15.8 s and 15.4 s) in both injection windows. The absence of a predictive model means LQI can only respond reactively once the error has accumulated in the integral state, resulting in a larger initial excursion and a slower return to steady state.

SC-MPC demonstrates the strongest local rejection capability of all controllers, recovering from both disturbances faster than any other strategy (10.3 s and 10.7 s recovery times) and recording the lowest segmented IAE in both injection windows. The scenario-based robust optimisation maintains feasibility margins around the current operating point even under perturbation, providing the optimiser with sufficient room to apply a strong corrective action without violating the output bounds. This advantage is structural: by enforcing constraints across multiple uncertainty realisations simultaneously, SC-MPC implicitly prepares a correction that is robust to the uncertainty about how the disturbance will propagate through the coupled dynamics.

SL-MPC produces the largest peak deviation on the h_2 injection (2.957 cm), an unexpected result given its strong tracking performance. The re-linearisation lag following the abrupt state change introduced by the injection explains this behaviour: the Jacobian update requires one or two samples to stabilise at the new perturbed operating point, during which the prediction model is temporarily less accurate than a fixed model that does not need to re-adapt. This limitation is specific to impulsive disturbances and does not affect SL-MPC's performance under the smoother reference changes of the tracking tests.

Overall Controller Ranking No single controller dominates all metrics across all tests. The relative merits of each strategy are summarised below.

LQI provides the fastest transient response and the lowest global tracking error in symmetric tests, with minimal computational cost and straightforward offline tuning. Its limitations are the absence of explicit constraint handling, which produces the largest overshoot in large-step tests, and the reactive disturbance rejection, which results in the longest recovery times. LQI is the appropriate choice when tracking speed is the primary objective, operating conditions are well characterised, and disturbances are infrequent or small.

Nominal MPC adds explicit constraint enforcement and a finite-horizon cost structure that suppresses overshoot and reduces actuator activity relative to LQI. Its primary limitation is the fixed linearisation, which causes measurable tracking degradation in wide-range tests and under opposing reference demands. It represents the natural first step beyond LQI for applications where constraint handling is required but nonlinear operating range variations are limited.

SL-MPC extends Nominal MPC with online re-linearisation, providing a tangible advantage whenever the operating range is wide or the two channels operate at significantly different levels. It delivers the smoothest control signals of all controllers across every test and the best h_4 tracking in multi-step and differential scenarios. The re-linearisation lag under impulsive disturbances is its only identified limitation. SL-MPC is the recommended choice when smooth actuation and nonlinear accuracy are simultaneously required without the overhead of explicit uncertainty quantification.

SC-MPC provides the most energy-efficient operation and the fastest disturbance rejection of all controllers, at the cost of conservative settling during large transients. The scenario-based tightening guarantees constraint satisfaction under the assumed parametric uncertainty with 95% probability, making it the appropriate choice when safety margins must be maintained under valve coefficient uncertainty without online identification.

Adaptive SC-MPC combines the constraint robustness of SC-MPC with online learning that progressively reduces conservatism as data accumulates. It improves on SC-MPC across all key metrics in the setpoint tracking test, with the improvement growing as the parameter box contracts. It represents the most capable controller in the comparison but also the most complex, requiring careful tuning of the identification safety mechanisms (guard period, cooldown, derivative threshold) to ensure stable real-plant operation. It is the appropriate choice when parameters drift over time or differ from their identified values, and when the operational duration is sufficient for the set membership box to contract to a useful degree.

6 | Conclusions and Future Work

This thesis presented a systematic experimental study of five control strategies of increasing complexity applied to the laboratory quadruple-tank process: Linear Quadratic Regulator with integral action (LQI), nominal velocity-form MPC, MPC with successive linearisation (SL-MPC), scenario-based MPC (SC-MPC), and adaptive scenario-based MPC with zonotopic set membership identification (Adaptive SC-MPC). All controllers were designed from the same grey-box model and validated on the physical Quanser plant under a common experimental protocol, enabling a rigorous quantitative comparison across seven test scenarios that stress initial filling transients, multi-step and asymmetric reference tracking, near-saturation large-signal operation, and unmeasured disturbance rejection.

6.1. Main Findings

LQI achieves the fastest transient response and the lowest global IAE across all symmetric tracking tests, settling h_2 in 25.8 s in Test 1 compared to 35.5 s for Nominal MPC and 49.0 s for SC-MPC. This advantage is purchased through aggressive integral action that saturates the pump voltage at 18 V immediately after reference changes, producing the highest overshoot of all configurations (13.22% on h_2 in Test 5) and the highest energy consumption in every test. The absence of predictive constraint handling results in the poorest disturbance rejection: LQI records the largest peak deviations (2.031 cm on h_2) and the slowest recovery times (15.8 s) in the injection test, confirming that reactive integral action cannot exploit knowledge of the coupled tank dynamics to anticipate disturbance propagation. Anti-windup protection proved essential for stable operation during the initial filling phase from empty.

Nominal MPC enforces explicit level and voltage bounds through the QP optimisation and reduces overshoot to below 6% in Test 1 relative to LQI, confirming the value of predictive constraint handling even at a fixed linearisation. However, the fixed Jacobian at the 10 cm nominal operating point is a measurable limitation whenever the system operates far from this design centre: the settling time on h_4 degrades to 39.8 s in Test 5, the slowest of any controller on any channel in the maximum-limit test, and the h_2 return overshoot reaches 34.53% under opposing reference demands in Test 4. These results confirm that nominal linear MPC is adequate for small-range regulation but requires augmentation to handle the full operating envelope of the quadruple-tank process, as also observed in comparable benchmark studies [22], [25].

SL-MPC resolves the operating-point sensitivity of Nominal MPC through online re-linearisation at each sampling instant. The updated Jacobian tracks the level-dependent Torricelli gain continuously, providing a measurable advantage on h_4 tracking across all multi-step tests: the lowest h_4 IAE of 55.8 cm s in Test 2, the fastest h_4 settling of 24.4 s during the initial fill, and the best differential tracking performance in Test 4 under opposing reference demands. SL-MPC also delivers the smoothest control signals of all strategies in every test, with total variation consistently 10–15% below Nominal MPC and SC-MPC, in agreement with the smoothness property reported for successive linearisation approaches in the literature [13], [14]. The one identified limitation is the re-linearisation lag following impulsive disturbances: the Jacobian requires one to two samples to re-stabilise after the abrupt state change introduced by the manual injection, producing the largest peak deviation on the h_2 injection (2.957 cm) in the disturbance rejection test.

SC-MPC provides the most energy-efficient operation of all controllers, achieving the lowest total energy in every test (14 954 V^2 s in Test 1 to 20 292 V^2 s in Test 5) and the fastest local disturbance rejection (10.3 s and 10.7 s recovery times). The scenario-based constraint tightening with $K = 39$ scenarios at a 5% risk level [30], [31] maintains the effective level bounds $[2.19, 24.91]$ cm without a single violation across all seven tests, including the near-saturation operation of Test 5 and the manual injection test, providing strong empirical support for the probabilistic constraint satisfaction guarantee [8]. The cost of this robustness is the most conservative transient behaviour: SC-MPC consistently shows the longest settling times during large initial steps, as the scenario margins limit the control authority available when the operating point is far from steady state.

Adaptive SC-MPC demonstrates that the conservatism of the fixed scenario approach can be progressively reduced through online set membership identification of the valve coefficients [10], [33]. Under identical experimental conditions, the adaptive variant improves on SC-MPC by 15% on h_2 IAE, 26% on h_2 overshoot and 4% on total energy, as the parameter box contracts from the initial range $[0.75, 0.90]\,A^{\mathrm{id}}$ toward the true valve values. The scenario envelope width reduces from approximately 0.9 cm to 0.25 cm within the first 60 s, confirming that online learning provides a quantifiable reduction in conservatism on the real plant consistent with the framework of Shi et al. [36], [37]. The safety mechanisms implemented for real-plant operation — the 40-sample guard period, the 100-sample reference-change cooldown, and the 0.3 cm/s derivative threshold on the level measurements — proved essential: disabling the cooldown during pulse tracking caused controller instability, while the full mechanism maintained stable adaptive operation throughout the campaign. All five controllers completed their computations within the $T_s = 1\,\mathrm{s}$ sampling period, with execution times ranging from below 1 ms for LQI to approximately 60 ms for Adaptive SC-MPC, confirming feasibility on standard laboratory hardware.

6.2. Future Work

Several directions emerge directly from the limitations and open questions identified during this work.

Black-box ARX identification for the adaptive SC-MPC. The set membership identification implemented in this thesis operates directly on the physical mass balance equations, using the four valve outlet areas $A_1, \ldots, A_4$ as uncertain parameters. This grey-box approach requires knowledge of the tank cross-section, pump flow gains and split ratios, all of which were identified offline before the experiments. An alternative is to replace the physical regressor with a black-box AutoRegressive with eXogenous input (ARX) model, as proposed by Meza et al. [33], [35] and extended to zonotopic representations by Shi et al. [37]. In this formulation, the set membership identification updates the ARX parameter vector directly from input-output data, without requiring any structural knowledge of the plant. For the quadruple-tank system, where the pump flow maps and split ratios drift over time due to wear and valve fouling, an ARX-based adaptive SC-MPC would be more robust to plant-model mismatch beyond the valve coefficient space and could capture unmodelled nonlinearities in the pump characteristics as part of the identified input coefficients. The initial identification phase proposed in [33] — a short unconstrained excitation period before closed-loop operation begins — would be particularly natural to implement on the Quanser platform, where the tanks can be pre-filled before the reference is applied.

Improved sensor calibration and noise reduction on Tank 4. Throughout the experimental campaign, the Tank 4 pressure sensor exhibited a higher noise floor than the Tank 2 sensor, which contributed to the conservative choice of the noise bound $\varepsilon_{\mathrm{smi}}$ in the set membership update and to the larger steady-state scatter observed on h_4 across all controllers. A dedicated sensor recalibration campaign for Tank 4, combined with a low-pass pre-filter tuned to the $T_s = 1\,\mathrm{s}$ sampling period, would reduce the effective noise level seen by the identification algorithm and allow a tighter noise bound to be used. This would directly accelerate the contraction of the parameter box during the initial learning phase, since the strip width in (4.45) scales with ε_{smi}. The improved sensor signal would also reduce the likelihood of degenerate strip intersections, which currently trigger the rejection safeguard and suspend learning, and would provide a cleaner derivative estimate for the cooldown trigger threshold.

Full simulation validation of the Adaptive SC-MPC across all test scenarios. The experimental validation of the adaptive SC-MPC in this thesis is limited to Test 1 (setpoint tracking), with the comparative metrics reported in Table 5.19. The simulation environment developed in MATLAB and Simulink provides a deterministic nonlinear model of the quadruple-tank plant that is free from sensor noise and manual injection variability, making it the appropriate platform to validate the adaptive controller across the full test suite used for the other four strategies. In particular, multi-step tracking (Tests 2 and 3) would allow the learning dynamics to be studied across multiple reference changes, where the parameter box must contract and then remain stable across the cooldown phases triggered by each step. The differential test (Test 4) would expose the asymmetric identification behaviour that arises when h_2 and h_4 operate at significantly different levels simultaneously, stressing the channel-independence assumption of the interval-hull representation. Completing this simulation suite would provide a controlled baseline against which the real-plant results can be compared and would allow the safety mechanisms to be tuned more systematically before hardware deployment.

Full zonotopic uncertainty representation. The current implementation uses an axis-aligned box (interval hull) to represent the feasible parameter set, chosen because the four valve coefficients appear independently in the tank mass balances and because the box update is computationally trivial at $T_s = 1\,\text{s}$. A full first-order zonotopic representation, as formalised by Alamo et al. [12] and extended to constrained zonotopes by Scott et al. [15], would capture cross-parameter correlations that arise indirectly through the shared pump flow and split-ratio measurements. The zonotopic intersection can be computed as a sequence of strip intersections with predictable computational cost [37], and recent implementations on battery energy storage systems demonstrate that the overhead remains compatible with sampling periods shorter than 1 s [38]. For the quadruple-tank system, the expected benefit of zonotopic over box representation is modest when the valve coefficients are truly independent, but it would complete the theoretical connection to the framework of Meza et al. [33] and provide a tighter uncertainty set under operating conditions where the pump maps introduce correlated uncertainty across channels.

Active excitation and dual control. The set membership identification used in this thesis is entirely passive: the controller minimises the tracking cost and the identification simply processes the resulting closed-loop data. During the reference-change cooldown phases, learning is suspended and the uncertainty set does not contract. An active excitation strategy, in which a small persistent dither signal is superimposed on the optimal control input during steady-state periods, would maintain informative data even when the reference is constant and accelerate box contraction. The dual control framework, discussed in the context of scenario-based adaptive MPC by Meza et al. [33], formalises this trade-off by including an exploration term in the MPC cost that explicitly penalises uncertainty volume. Implementing a lightweight version of this idea on the quadruple-tank plant — for example, a scheduled low-amplitude pump voltage perturbation applied only when the derivative threshold confirms steady state — would be a natural extension of the cooldown mechanism already present in the current implementation and would reduce the dependence on reference changes as the primary source of identification data.

Offset-free tracking with explicit disturbance estimation. The velocity-form MPC formulation adopted throughout this thesis achieves offset-free tracking structurally, without an explicit disturbance observer. However, the persistent steady-state error of approximately 0.35 cm observed on h_4 in Test 4 under all three MPC controllers indicates that the implicit integral action is insufficient to fully reject the residual coupling disturbance introduced by the asymmetric flow routing. Augmenting the prediction model with an integrating disturbance state estimated by a Kalman filter, following the offset-free MPC framework of Pannocchia et al. [5], [41] and the robust formulation of Betti et al. [4], would provide an explicit rejection mechanism for this class of persistent cross-coupling disturbance. This extension is particularly relevant for the adaptive SC-MPC, where the parameter learning and the disturbance estimation could be designed jointly to distinguish between valve drift (captured by the set membership box) and unmodelled coupling effects (captured by the disturbance state).

Bibliography

[1] M. Farina, R. Scattolini, L. Fagiano, and A. La Bella, *Model predictive control: Ph.D. course lecture series*, Dipartimento di Elettronica, Informazione e Bioingegneria, 2024.

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

[3] D. Q. Mayne, "Model predictive control: Recent developments and future promise," *Automatica*, 2014.

[4] G. Betti, M. Farina, and R. Scattolini, "A robust MPC algorithm for offset-free tracking of constant reference signals," *IEEE Transactions on Automatic Control*, vol. 58, no. 9, pp. 2394–2400, 2013.

[5] G. Pannocchia, M. Gabiccini, and A. Artoni, "Offset-free MPC explained: Novelties, subtleties, and applications," *IFAC-PapersOnLine*, vol. 48, no. 23, pp. 342–351, 2015.

[6] I. A. Shehu and N. Abdul Wahab, "Applications of mpc and pi controls for liquid level control in coupled-tank systems," in *2016 IEEE International Conference on Automatic Control and Intelligent Systems (I2CACIS)*, IEEE, 2016, pp. 119–124.

[7] 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.

[8] G. Schildbach, L. Fagiano, C. Frei, and M. Morari, "The scenario approach for stochastic model predictive control with bounds on closed-loop constraint violations," *Automatica*, vol. 50, no. 12, pp. 3009–3018, 2014.

[9] M. Bauer, A. Horch, L. Xie, M. Jelali, and N. F. Thornhill, "The current state of control loop performance monitoring," *Journal of Process Control*, 2016.

[10] M. Tanaskovic, L. Fagiano, R. Smith, and M. Morari, "Adaptive receding horizon control for constrained MIMO systems," *Automatica*, vol. 50, no. 12, pp. 3019–3029, 2014.

[11] M. Milanese and A. Vicino, "Optimal estimation theory for dynamic systems with set membership uncertainty: An overview," *Automatica*, vol. 27, no. 6, pp. 997–1009, 1991.

[12] T. Alamo, J. M. Bravo, and E. Camacho, "Guaranteed state estimation by zonotopes," *Automatica*, vol. 41, no. 6, pp. 1035–1043, 2005.

[13] A. Zhakatayev, B. Rakhim, O. Adiyatov, A. Baimyshev, and H. A. Varol, "Successive linearization based model predictive control of variable stiffness actuated robots," in *IEEE International Conference on Advanced Intelligent Mechatronics (AIM)*, 2017.

[14] J. Berberich, J. Köhler, M. Müller, and F. Allgöwer, "Linear tracking MPC for nonlinear systems part i: The model-based case," *IEEE Transactions on Automatic Control*, vol. 67, no. 9, pp. 4426–4441, 2022.

[15] J. K. Scott, D. M. Raimondo, G. R. Marseglia, and R. D. Braatz, "Constrained zonotopes: A new tool for set-based estimation and fault detection," *Automatica*, vol. 69, pp. 126–136, 2016.

[16] F. Bonassi, M. Farina, J. Xie, and R. Scattolini, "On recurrent neural networks for learning-based control: Recent results and ideas for future developments," *Journal of Process Control*, 2024.

[17] L. Boca de Giuli, A. La Bella, G. De Nicolao, and R. Scattolini, "Physics-informed neural network modeling and predictive control of district heating systems," *IEEE Transactions on Control Systems Technology*, vol. 32, no. 4, 2024.

[18] M. Lévis and T. Lee, "The quanser platform for control systems research validation," Quanser, White Paper, 2013.

[19] E. Calle and J. Oliden, "Recurrent neural network based predictive control applied to a four coupled tank system," *Journal of Process Control*, 2023.

[20] I. Alvarado et al., "A comparative analysis of distributed MPC techniques applied to the HD-MPC four-tank benchmark," *Journal of Process Control*, vol. 21, no. 6, pp. 800–815, 2011.

[21] M. Mercangöz and F. J. Doyle, "Distributed model predictive control of an experimental four-tank system," *Journal of Process Control*, 2007.

[22] P. C. Blaud, P. Chevrel, F. Claveau, P. Haurant, and A. Mouraud, "Four MPC implementations compared on the quadruple tank process benchmark: Pros and cons of neural MPC," *IFAC-PapersOnLine*, vol. 55, no. 16, pp. 344–349, 2022.

[23] M. M. Gulzar, M. Munawar, Z. Dewan, M. Salman, and S. Iqbal, "Level control of coupled conical tank system using adaptive model predictive controller," *IEEE Access*, 2019.

[24] H. A. F. Almurib, M. Askari, and M. Moghavvemi, "Model predictive control of quadruple tanks system," *International Journal of Control and Automation*, 2011.

[25] A. H. D. Christensen, T. K. S. Ritschel, J. L. Svensen, S. Hørsholt, J. K. Huusom, and J. B. Jørgensen, "Comparing model-based control strategies for a quadruple tank system: Decentralized PID, LMPC and NMPC," *Journal of Process Control*, 2023.

[26] E. Rosman, S. N. Mohd Azam, and M. Che Razali, *Control strategies for the quadruple tank system: A comparative study of PID and MPC*, Unpublished or conference paper PDF, Provided as reference material.

[27] S. Özkan, T. Kara, and M. Aras, *Modelling, simulation and control of quadruple tank process*, Unpublished or conference paper PDF, Provided as course or reference material.

[28] M. Askari, M. Moghavvemi, H. A. F. Almurib, and K. M. Muttaqi, "Multivariable offset-free model predictive control for quadruple tanks system," *IEEE Transactions on Industry Applications*, vol. 52, no. 2, pp. 1882–1890, 2016.

[29] M. Murillo, G. Sánchez, and L. Giovanini, "Iterated non-linear model predictive control based on tubes and contractive constraints," *ISA Transactions*, vol. 62, pp. 120–128, 2016. DOI: `10.1016/j.isatra.2016.01.008`.

[30] M. C. Campi and S. Garatti, "The exact feasibility of randomized solutions of uncertain convex programs," *SIAM Journal on Optimization*, vol. 19, no. 3, pp. 1211–1230, 2008.

[31] G. C. Calafiore and M. C. Campi, "The scenario approach to robust control design," *IEEE Transactions on Automatic Control*, vol. 51, no. 5, pp. 742–753, 2006.

[32] H. S. Abbas, R. Tóth, M. Petreczky, N. Meskin, J. Mohammadpour Velni, and P. Koelewijn, "LPV modeling of nonlinear systems: A multi-path feedback linearization approach," *International Journal of Robust and Nonlinear Control*, vol. 31, no. 18, pp. 9436–9465, 2021.

[33] E. Meza, A. Del Duca, F. Ruiz, and C. Ocampo-Martinez, "Adaptive scenario-based predictive control: A set membership learning approach," *IFAC-PapersOnLine*, vol. 59, no. 26, pp. 323–328, 2025.

[34] A. Cordoba-Pacheco and F. Ruiz, "Optimal energy management in multi-microgrids: A scenario-based MPC approach," *IEEE Transactions on Smart Grid*, 2023.

[35] E. M. Meza Camargo, "Scenario-based MPC with adaptive ARX set membership identification," M.S. thesis, Universitat Politècnica de Catalunya, 2024.

[36] Q. Shi, A. Cordoba-Pacheco, and F. Ruiz, "Zonotopic uncertainty based adaptive scenario model predictive control: A comparative study on BESS systems," *IEEE Transactions on Control Systems Technology*, 2025.

[37] Q. Shi, A. Cordoba-Pacheco, and F. Ruiz, "Efficient learning control through zonotopic set membership estimation and scenario MPC," *IFAC-PapersOnLine*, 2024.

[38] A. Cordoba-Pacheco, L. Petrucci, M. Filippone, and F. Ruiz, "Experimental data-driven BESS controller tuning: A set membership approach," in *IEEE Conference on Control Technology and Applications (CCTA)*, 2024.

[39] W. D'Amico, A. La Bella, and M. Farina, "An incremental input-to-state stability condition for a class of recurrent neural networks," *IEEE Transactions on Automatic Control*, vol. 69, no. 4, 2024.

[40] A. Cordoba Pacheco, *Advanced and multivariable control - AMC tutoring: Lecture 1 - identification*, Lecture Notes, Politecnico di Milano, Course material for AMC Tutoring sessions, 2024.

[41] G. Pannocchia, "Offset-free tracking MPC: A tutorial review and comparison of different formulations," in *European Control Conference (ECC)*, 2015.

Acknowledgements

I would like to express my sincere gratitude to my supervisor, Prof. Fredy Orlando Ruiz Palacios, for providing the validated grey-box model and identification data that served as the foundation for this research, and for his guidance throughout this work.

I am deeply thankful to my co-advisors, Andres Felipe Cordoba Pacheco and Eva Masero Rubio, for their continuous support, insightful discussions, and valuable feedback during the development of this thesis.

I would also like to thank the laboratory staff at Politecnico di Milano for their assistance with the experimental setup and for ensuring the safe and reliable operation of the quadruple tank system.

Finally, I am grateful to my family and friends for their unwavering support and encouragement throughout my master's studies.