energies

MDPI

Article

Inlet-Temperature Disturbance Rejection in Solar Collector Fields Using Learning-Based Practical Nonlinear Predictive Control

Igor M. L. Pataro [1,*], Juan D. Gil [2], José L. Guzmán [2], Manuel Berenguel [2] and Teodoro Alamo [1]

[1] Departamento de Ingeniería de Sistemas y Automática, Universidad de Sevilla, 41092 Seville, Spain; talamo@us.es
[2] Centro Mixto CIESOL, ceiA3, Universidad de Almería, Ctra. Sacramento s/n, 04120 Almería, Spain; juandiego.gil@ual.es (J.D.G.); joguzman@ual.es (J.L.G.); beren@ual.es (M.B.)
* Correspondence: imendeslima@us.es

Abstract

Rapid variations in the Heat-Transfer-Fluid (HTF) inlet temperature can disrupt outlet-temperature regulation in solar collector fields, particularly when fast recovery interacts with delayed control action. This work assesses the previously developed Learning-Based Practical Nonlinear Model Predictive Controller (LBPNMPC) for rejecting such disturbances, motivated by expansion-vessel fluid-exchange transients observed at an actual facility. A matched clear-day simulation uses fourteen inlet-temperature disturbances generated from five experimentally measured profiles. LBPNMPC is compared with conventional filtered-error PNMPC, both with measured-disturbance feedforward, and with PNMPC without disturbance feedforward. In this matched simulation, LBPNMPC reduces the inlet-event integral absolute error by 34.3%, integral squared error by 49.3%, mean recovery time by 23.9%, and HTF-flow total variation by 11.4% relative to conventional filtered-error PNMPC. Experimental operation is assessed at the CIESOL research centre, University of Almería, Spain. During a severe 9.27 °C inlet-temperature decrease, the corresponding measured outlet-temperature decrease was 5.60 °C, approximately 60% of the inlet-temperature variation. The experiment demonstrates the practical applicability of LBPNMPC during a severe inlet-temperature disturbance, complementing the quantitative comparison with conventional PNMPC provided by the matched simulation.

Keywords: model predictive control; nonlinear predictive control; disturbance rejection; solar collector field; learning-based control; solar thermal energy

Check for updates

Academic Editor: Tapas Mallick

Received: 14 October 2026
Revised: 14 November 2026
Accepted: 17 November 2026
Published: 19 November 2026

Copyright: © 2026 by the authors. Licensee MDPI, Basel, Switzerland. This article is an open access article distributed under the terms and conditions of the Creative Commons Attribution (CC BY) license.

1. Introduction

The main control objective in solar thermal plants is to regulate the heat-transfer-fluid (HTF) outlet temperature while efficiently exploiting the available solar energy. Since solar irradiance cannot be manipulated, the outlet temperature is commonly controlled by adjusting the HTF flow rate according to the available thermal power and the plant operating conditions [1,2]. Several disturbances affect this regulation, including solar irradiance, ambient conditions, optical efficiency, and HTF inlet temperature [3]. Among them, the inlet temperature directly determines the thermal state of the fluid entering the solar collector field (SCF) and, therefore, constitutes an important measured disturbance for temperature control.

Under normal operating conditions, the inlet temperature generally evolves more slowly than solar irradiance and can reasonably be assumed constant over a short prediction horizon [4]. However, hydraulic events can cause much faster variations. In particular, fluid

exchange between the expansion vessel and the main hydraulic circuit can introduce HTF at a different thermal level, causing a sudden decrease in the inlet temperature followed by rapid recovery. This transient is particularly demanding from a control perspective because the decrease in inlet temperature induces a compensating variation in the HTF flow, while both the disturbance and manipulated input propagate through the SCF with transport delays. Consequently, the inlet temperature may recover while the preceding control action is still propagating through the field, potentially producing a significant outlet-temperature excursion.

This phenomenon has been recognised in solar thermal field operation. In ref. [5], inlet-temperature transients were considered within a hierarchical control framework for the start-up procedure of a solar thermal field with direct storage. In that work, a decision-making layer determined whether the control loops should use the measured inlet temperature or a process-model-estimated temperature. When a large transient was detected, the measured value was temporarily replaced by the model-based estimate, thereby preventing the anomalous measurement from directly affecting the corresponding control loops.

Herein, rather than replacing the measured inlet temperature during rapid transients, the actual inlet-temperature measurement is retained and explicitly incorporated into the predictive controller as a measured disturbance. Model Predictive Control (MPC) is particularly suitable for this problem because it can explicitly incorporate measured disturbances, dynamic predictions, and actuator constraints into the control calculation [6–8]. For nonlinear solar thermal processes, Practical Nonlinear Model Predictive Control (PNMPC) provides a computationally efficient alternative to conventional nonlinear MPC by repeatedly linearising a nonlinear plant model and solving a quadratic optimisation problem [9]. PNMPC has been successfully applied to SCF temperature regulation, providing satisfactory control performance with relatively low online computational requirements [10,11]. Nevertheless, the resulting predictions may contain error due to remaining plant–model mismatch, transport-delay uncertainty, and unmodelled process dynamics.

Other predictive-control approaches have also addressed disturbance rejection in SCFs [12]. Since MPC relies on model predictions to compute control actions, plant–model mismatch can reduce prediction accuracy, introduce steady-state tracking offsets, and deteriorate disturbance-rejection performance if not adequately compensated [13,14]. To address these challenges, different disturbance-compensation strategies have been investigated [15]. In ref. [16], measured-disturbance feedforward was incorporated into a gain-scheduled predictive-control strategy, accounting for inlet-temperature dynamics and the impact of future disturbance information. In ref. [17], Kalman-filter-based state estimation was combined with predictive control to improve robustness against unmeasured disturbances. More recently, ref. [18] proposed a frequency-domain disturbance-estimation approach within MPC, particularly for periodic disturbances. In ref. [19], the authors combined lumped-disturbance estimation and prediction with constrained predictive optimisation for a parabolic trough solar field. Their approach uses an extended state observer to estimate the lumped disturbance and a least-squares support vector machine to predict its future evolution. Zhang et al. [20] investigated stochastic dual MPC for solar thermal collector fields, combining online parameter estimation through Kalman filtering with probabilistic output constraints. Their formulation incorporates an exploration mechanism to improve parameter learning while regulating the outlet temperature. Beyond solar applications, in [21] the authors combined generalised predictive control with generalised proportional integral (GPI) observers to estimate and compensate for lumped disturbances in irrigation canals. Despite these developments, the combined effects of abrupt inlet-temperature variations, flow-dependent transport delays, and residual model-prediction errors remain

challenging. The present work, therefore, investigates adaptive prediction-error correction to improve disturbance rejection under these conditions.

The PNMPC framework has considered different prediction-error compensation mechanisms. An integral correction based on the filtered prediction error was investigated in [10], while a filtered-error formulation aimed at improving robustness to modelling inaccuracies was subsequently studied in [11]. More recently, the Learning-Based Practical Nonlinear Model Predictive Controller (LBPNMPC) introduced in [4] replaced this fixed correction with a parametric Oracle function. More broadly, learning-based MPC seeks to improve model-based predictive control by incorporating information learned from process data, allowing the prediction model or control strategy to adapt to uncertainties while retaining the ability to account for operating constraints [22,23]. In the case of the LBPNMPC, the parametric Oracle estimates the future evolution of the residual prediction error from past plant information, using Recursive Least Squares with Exponential Resetting (ER + RLS) to update its parameters online [24] and a multistep prediction structure related to adaptive predictive control [25]. This principle is applied specifically to the residual prediction error, preserving the original nonlinear plant model while adaptively correcting its predictions.

Although the previous prediction-error compensation strategies improve robustness to modelling inaccuracies, their predefined correction dynamics may be less effective when abrupt inlet-temperature variations generate rapidly changing model–process discrepancies. In this context, an adaptive prediction-error correction that can adjust to transient behaviour is particularly relevant to ensure effective disturbance rejection. Therefore, this work applies the adaptive LBPNMPC to improve rejection of rapid inlet-temperature transients caused by hydraulic events in SCFs and to mitigate the resulting outlet-temperature variations. The hypothesis is that retaining the measured inlet temperature while adaptively correcting the residual model error improves rejection of rapid inlet-temperature disturbances, particularly during the recovery phase, when the disturbance interacts with delayed control action. In this work, both conventional PNMPC and LBPNMPC are tested and compared using the same nonlinear model, measured disturbances, disturbance-feedforward contribution, transport delays, constraints, and optimisation problem. Their main difference is how they treat the residual prediction error: conventional PNMPC applies a predefined filtered-error correction, whereas the LBPNMPC Oracle generates an adaptive multistep prediction from recent plant data.

The LBPNMPC formulation is retained from [4]. The scientific contribution of the present work is a focused assessment of rapid inlet-temperature disturbances and their interaction with delayed flow compensation. This assessment combines a measured-profile benchmark, comparisons of feedforward and implemented residual-error compensation, and joint interpretation of event-based errors, complete-record performance, recovery time, and control activity. The main contributions are summarised as follows:

1. Measured inlet-temperature transients associated with operator-observed expansion-vessel valve opening are characterised using CIESOL operating data, with approximate thermal-transient intervals identified retrospectively from the temperature record.
2. A simulation benchmark constructed from five measured profiles introduces fourteen inlet-temperature valleys with prescribed amplitudes and timing to assess disturbance rejection under the inherited delay-mismatch scenario.
3. Adaptive and filtered prediction-error-compensation schemes are compared under common plant and MPC settings and imposed inlet-temperature profiles. The comparison includes LBPNMPC's excitation for Oracle estimation, while a PNMPC-noFF case assesses the measured-disturbance feedforward contribution. The results also

demonstrate why local error integrals require interpretation alongside recovery and complete-record performance.

4. A CIESOL implementation demonstrates practical controller operation during inlet-temperature transients, complementing the simulation assessment without establishing experimental comparative superiority.

This work is structured as follows. Section 2 describes the CIESOL SCF, its nonlinear model, and the inlet-temperature disturbance considered in this study. Section 3 presents the predictive-control formulations and their treatment of the measured inlet-temperature disturbance. Sections 4 and 5 report the simulation and experimental assessments, respectively. Section 6 discusses the disturbance-rejection mechanisms, practical implications, and limitations of the study, while Section 7 summarises the main conclusions.

2. Solar Collector Field and Disturbance-Rejection Problem

The inlet-temperature phenomenon introduced in Section 1 is analysed using the SCF of the CIESOL research centre at the University of Almería, Spain, as a case study. CIESOL is an experimental facility dedicated to research on solar-energy technologies, including solar thermal generation, energy conversion, and advanced optimisation, modelling and control solutions. The main goal of the solar thermal plant is to provide thermal energy to the building's indoor climate system, heating or cooling the indoor rooms.

The plant consists of 80 Solaris CP1 flat-plate solar collectors. Each collector has an area of approximately 2 m^2, providing a total collecting area of approximately 160.2 m^2. The collectors are oriented south and tilted 20° above the horizontal. Hydraulically, they are arranged in 10 parallel rows, with eight collectors connected in series within each row. This configuration allows the facility to operate over a broad range of outlet temperatures and HTF flow rates, making it suitable for experimentally assessing advanced control strategies. Figure 1a shows an overview of the facility, while Figure 1b presents the SCF considered in this work.



(**a**) (**b**)

Figure 1. CIESOL solar-energy research facility and solar collector field: (**a**) overview of the CIESOL building at the University of Almería; (**b**) solar collector field considered in this work.

The facility has previously been used to model and experimentally assess predictive-control strategies [4]. In this work, the CIESOL installation is particularly useful because its hydraulic configuration includes an expansion vessel connected to the primary circuit near the circulation pump and the collector-field inlet, allowing inlet-temperature transients to be observed during normal operation.

2.1. CIESOL Solar Collector Field and Model

Figure 2 presents a simplified schematic of the CIESOL facility, including the main components of the hydraulic circuit relevant to this study. For clarity, the simulation omits the gas heater, as it represents the total thermal power provided by the SCF. The primary

circuit circulates the HTF through the solar collector field and exchanges the collected thermal energy with the storage and secondary circuits. The HTF flow rate, q, is manipulated through the circulation pump, while the collector outlet temperature, T_{out}, is the controlled variable. The main measured disturbances are the collector inlet temperature, T_{in}, solar irradiance, I, and ambient temperature, T_a.

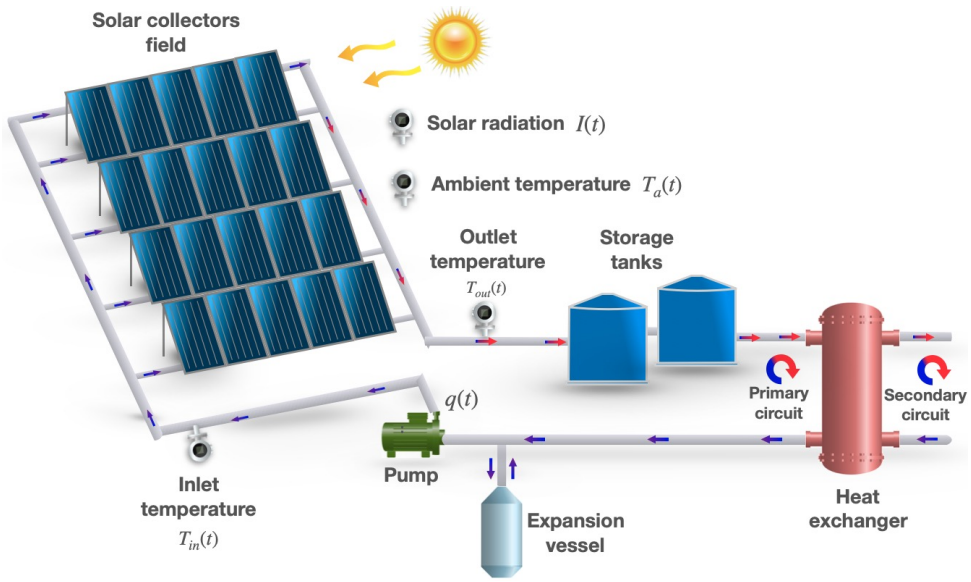


Figure 2. Simplified scheme of the CIESOL solar thermal facility.

The lumped-parameter model for the CIESOL SCF thermal dynamics was previously identified and validated in [4], based on the energy-balance formulation commonly used for solar collector fields. The outlet-temperature dynamics are given by

$$\begin{aligned}\frac{\mathrm{d}T_{\text{out}}(t)}{\mathrm{d}t} = {} & \frac{\beta}{\rho C_p A_{sc}} I(t) - \frac{H}{\rho C_p A_{sc} L}\left[\frac{T_{\text{out}}(t) + T_{\text{in}}(t - t_{d,T})}{2} - T_a(t)\right] \\ & - \frac{q(t - t_{d,q})}{A_{sc} c_f L}\left[T_{\text{out}}(t) - T_{\text{in}}(t - t_{d,T})\right],\end{aligned} \tag{1}$$

where $T_{\text{out}}(t)$ and $T_{\text{in}}(t)$ are the outlet and inlet temperatures of the HTF, $T_a(t)$ is the ambient temperature, $I(t)$ is the solar irradiance, and $q(t)$ is the total HTF volumetric flow rate. Moreover, A_{sc} is the collector-pipe cross-sectional area, L is the equivalent absorber-tube length, C_p and ρ are the HTF specific heat capacity and density, respectively, β represents the irradiance efficiency factor, H is the global heat-loss coefficient, and c_f is the flow-unit conversion factor. The transport delays associated with the inlet temperature and HTF flow are denoted by $t_{d,T}$ and $t_{d,q}$, respectively, and they depend on the water flow rate. In this work, $t_{d,T}$ and $t_{d,q}$ are approximated as the mean value obtained empirically (and the uncertainty associated with this choice will be addressed by the proposed control structure).

The obtained values adopted in this study are those identified for the CIESOL facility in [4] and are summarised in Table 1.

The model parameters were previously identified from experimental data and validated over different operating temperatures and HTF flow rates, obtaining $R^2 = 0.9847$ and a mean prediction error of 0.594 °C. In this work, the validated model is used for simulation and controller prediction, focusing on the inlet-temperature disturbance rather than a new model-identification problem.

Table 1. Nominal parameters of the CIESOL SCF model.

Parameter	Description	Value	Unit
A_{sc}	Collector-pipe cross-sectional area	1.13×10^{-4}	m^2
c_f	Flow conversion factor	3.6×10^4	s/h
C_p	HTF specific heat capacity	4186	J/(kg °C)
H	Global heat-loss coefficient	102.52	W/°C
L	Equivalent absorber-tube length	46.6	m
β	Irradiance efficiency factor	0.3882	m
ρ	HTF density	972	kg/m^3
$t_{d,T}$	Inlet-temperature transport delay	41	s
$t_{d,q}$	HTF-flow transport delay	38	s

As highlighted in [26], the HTF transport delay depends on operating and geometric factors, including the flow rate, tube dimensions, and sensor location. Consequently, the delay t_{d_q} is not constant but varies with operating conditions and directly affects the accuracy of the predicted control response. This effect becomes particularly critical during the rapid inlet-temperature transients considered in this work. When T_{in} decreases, the controller compensates by reducing the manipulated flow $q(t)$; however, this flow reduction simultaneously modifies the HTF transport delay. Therefore, the disturbance not only alters the thermal state of the fluid entering the SCF, but also changes the timing with which the compensating control action propagates to the outlet. During the subsequent rapid recovery of T_{in}, this varying delay can amplify the mismatch between the predicted and actual plant response, thereby increasing the risk of outlet-temperature excursions. Section 4 analyses the LBPNMPC's ability to estimate the disturbance prediction effect.

2.2. Expansion-Vessel-Induced Inlet-Temperature Disturbance

The expansion vessel shown in Figure 2 is connected to the primary hydraulic circuit immediately upstream of the circulation pump. It accommodates pressure and volume variations of the HTF in the closed circuit. Fluid exchange between the expansion-vessel branch and the main circulation loop can consequently modify the thermal conditions at the SCF inlet.

Figure 3 presents a representative CIESOL operating period during which the on-site operator observed the opening of the valve connecting the expansion-vessel branch to the main hydraulic circuit. Only the occurrence of this opening was noted; its duration was not recorded, and the timing precision of the operator observation was not quantified. These observations were used to locate the relevant portions of the recorded inlet-temperature trajectory. The thermal-transient intervals were then identified retrospectively by visual inspection, from an abrupt departure of T_{in} from its local pre-event trend through its recovery towards normal evolution. The shaded regions, therefore, indicate approximate thermal-transient intervals rather than independently measured valve-opening or fluid-exchange durations. Since no dedicated vessel pressure or branch-flow measurement is available, the observations support an association with expansion-vessel exchange but do not independently establish its duration or exclude other contributing causes.

A physically plausible explanation for these transients is fluid exchange between the expansion vessel and the main hydraulic circuit, which can alter the thermal conditions at the SCF inlet. The direction of the resulting T_{in} variation depends on the temperature of the fluid exchanged with the vessel branch relative to that of the HTF circulating in the main loop. Because some fluid may remain in the expansion-vessel branch for a certain residence time and exchange heat with its surroundings, its temperature can differ from the instantaneous temperature of the circulating fluid. Consequently, the re-entry of cooler fluid

produces a valley in T_{in}, whereas the return of comparatively warmer fluid can generate a positive temperature peak. The magnitude and sign of these transients, therefore, depend on the hydraulic circuit's previous thermal history and on the direction and amount of fluid exchanged. Once the exchange ceases and normal circulation conditions are restored, T_{in} rapidly returns towards the temperature imposed by the main loop. This fast recovery, combined with the delayed effect of the control action and the thermal dynamics of the collector field, can then produce a significant excursion in T_{out}. The expansion vessel is not a manipulated element of the temperature-control system. Its effect enters the control problem through the measured inlet temperature. Hence, the controller must regulate T_{out} by manipulating q while the disturbance propagates from the inlet to the outlet with the transport delay $t_{d,T}$.

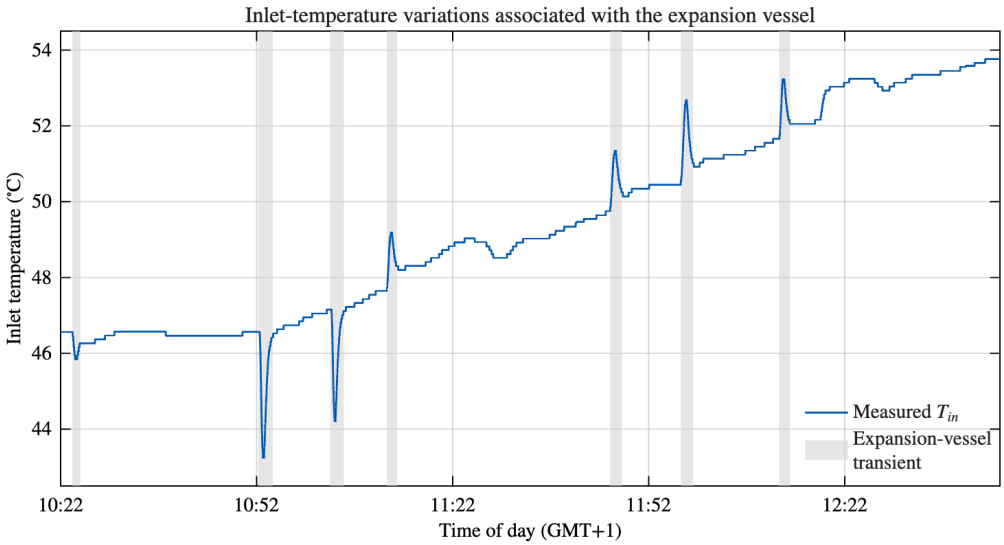


Figure 3. Representative CIESOL operating data showing inlet-temperature transients associated with operator-observed opening of the expansion-vessel connection valve. The shaded regions indicate approximate thermal-transient intervals identified retrospectively from the measured inlet-temperature trajectory. The operator observations indicate event occurrence only; valve-opening duration was not recorded, and the interval boundaries have unquantified timing uncertainty.

The rapid recovery of T_{in} is particularly important. From Equation (1), a decrease in inlet temperature tends to reduce T_{out}. The controller compensates for this effect mainly by decreasing q, thereby increasing the HTF's thermal gain as it crosses the collector field. If T_{in} subsequently recovers before the effect of the compensating flow action has fully propagated through the field, the delayed control action and recovered inlet temperature can act simultaneously, producing an outlet-temperature overshoot. This interaction between disturbance and control dynamics defines the disturbance-rejection problem investigated in the following sections.

2.3. Control Problem Formulation

The control objective is to regulate the SCF outlet temperature around the reference $r(k)$ while rejecting inlet-temperature transients and respecting the physical constraints imposed on the HTF flow. Here, k denotes the discrete-time index, with the corresponding continuous time given by $t = kT_s$, where T_s is the sampling period. The tracking error is defined as

$$e_y(k) = r(k) - T_{out}(k). \quad (2)$$

Because the principal interest of this study is the response to rapid inlet-temperature disturbances, controller performance is evaluated both over the complete operating interval and locally around each inlet-temperature event.

The disturbance-rejection performance is quantified using the Integral Absolute Error (IAE) and Integral Squared Error (ISE),

$$\mathrm{IAE} = T_s \sum_{k\in\mathcal{K}} |e_y(k)|, \qquad \mathrm{ISE} = T_s \sum_{k\in\mathcal{K}} e_y^2(k), \tag{3}$$

where $\mathcal{K}$ denotes either the complete analysed record or a response interval associated with an inlet-temperature disturbance. For event-based evaluation of the simulations in Section 4.2, $\mathcal{K}$ contains all samples within the corresponding inclusive window to evaluate the event metrics. Its boundaries are the zero-valued endpoints of the imposed inlet-temperature perturbation relative to its baseline. The total variation of the manipulated flow,

$$\mathrm{TV}(q) = \sum_{k=2}^{n} |q(k) - q(k-1)|, \tag{4}$$

is additionally used to evaluate the control activity required to reject the disturbances.

3. Predictive Control Strategies

This section summarises the PNMPC and LBPNMPC formulations used in the comparative analysis. Both controllers are based on the same nonlinear SCF model, prediction and control horizons, measured disturbances, constraints, and optimisation problem. Their main difference lies in how they compensate for the error between the nonlinear-model prediction and the measured plant response. The complete development of the PNMPC and LBPNMPC formulations can be found in [4,9].

3.1. Practical Nonlinear Model Predictive Control

Consider a nonlinear discrete-time model expressed in terms of past outputs, manipulated inputs, and measured disturbances as

$$\boldsymbol{y}(k+1) = f\left(\tilde{\boldsymbol{y}}(k), \tilde{\boldsymbol{u}}(k), \tilde{\boldsymbol{d}}(k)\right), \tag{5}$$

where $\boldsymbol{y} \in \mathbb{R}^{n_y}$, $\boldsymbol{u} \in \mathbb{R}^{n_u}$, and $\boldsymbol{d} \in \mathbb{R}^{n_d}$ denote the system outputs, manipulated inputs, and measured disturbances, respectively. The vectors $\tilde{\boldsymbol{y}}$, $\tilde{\boldsymbol{u}}$, and $\tilde{\boldsymbol{d}}$ collect the corresponding past values required by the nonlinear model.

In PNMPC, the nonlinear model is recursively simulated over the prediction horizon N_p to obtain the future output vector

$$\hat{\boldsymbol{Y}} = \begin{bmatrix} \hat{\boldsymbol{y}}(k+1|k) \\ \hat{\boldsymbol{y}}(k+2|k) \\ \vdots \\ \hat{\boldsymbol{y}}(k+N_p|k) \end{bmatrix}. \tag{6}$$

A first-order numerical linearisation is then performed around nominal future input and disturbance trajectories, denoted by $\boldsymbol{u}_0$ and $\boldsymbol{d}_0$. The resulting prediction can be written as

$$\hat{\boldsymbol{Y}} = \mathcal{F}\left(\tilde{\boldsymbol{Y}}, \tilde{\boldsymbol{u}}, \boldsymbol{u}_0, \tilde{\boldsymbol{d}}, \boldsymbol{d}_0\right) + \left.\frac{\partial \hat{\boldsymbol{Y}}}{\partial \Delta \boldsymbol{u}}\right|_{\boldsymbol{u}_0, \boldsymbol{d}_0} \Delta \boldsymbol{u} + \left.\frac{\partial \hat{\boldsymbol{Y}}}{\partial \Delta \boldsymbol{d}}\right|_{\boldsymbol{u}_0, \boldsymbol{d}_0} \Delta \boldsymbol{d}, \tag{7}$$

where $\mathcal{F}(\cdot)$ denotes the recursive simulation of the nonlinear model over N_p. The future increments are defined as

$$\boldsymbol{\Delta u} = \begin{bmatrix} \Delta \boldsymbol{u}(k|k) \\ \Delta \boldsymbol{u}(k+1|k) \\ \vdots \\ \Delta \boldsymbol{u}(k+N_c-1|k) \end{bmatrix}, \qquad \boldsymbol{\Delta d} = \begin{bmatrix} \Delta \boldsymbol{d}(k|k) \\ \Delta \boldsymbol{d}(k+1|k) \\ \vdots \\ \Delta \boldsymbol{d}(k+N_p-1|k) \end{bmatrix}, \tag{8}$$

with N_c denoting the control horizon.

When all future input and disturbance increments are set to zero, the nonlinear simulation gives the free response

$$\boldsymbol{F} = \mathcal{F}\left(\tilde{\boldsymbol{Y}}, \tilde{\boldsymbol{u}}, \boldsymbol{u}_0, \tilde{\boldsymbol{d}}, \boldsymbol{d}_0\right). \tag{9}$$

The forced responses associated with future input and disturbance variations are represented by the Jacobian matrices

$$\boldsymbol{G} = \left.\frac{\partial \hat{\boldsymbol{Y}}}{\partial \boldsymbol{\Delta u}}\right|_{\boldsymbol{u}_0, \boldsymbol{d}_0}, \qquad \boldsymbol{G}_d = \left.\frac{\partial \hat{\boldsymbol{Y}}}{\partial \boldsymbol{\Delta d}}\right|_{\boldsymbol{u}_0, \boldsymbol{d}_0}. \tag{10}$$

For a single-output system, the input forced-response matrix has the structure

$$\boldsymbol{G} = \begin{bmatrix} \frac{\partial \hat{y}(k+1|k)}{\partial \Delta u(k|k)} & 0 & \cdots & 0 \\ \frac{\partial \hat{y}(k+2|k)}{\partial \Delta u(k|k)} & \frac{\partial \hat{y}(k+2|k)}{\partial \Delta u(k+1|k)} & \cdots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ \frac{\partial \hat{y}(k+N_p|k)}{\partial \Delta u(k|k)} & \frac{\partial \hat{y}(k+N_p|k)}{\partial \Delta u(k+1|k)} & \cdots & \frac{\partial \hat{y}(k+N_p|k)}{\partial \Delta u(k+N_c-1|k)} \end{bmatrix}, \tag{11}$$

whereas the disturbance forced-response matrix is

$$\boldsymbol{G}_d = \begin{bmatrix} \frac{\partial \hat{y}(k+1|k)}{\partial \Delta \boldsymbol{d}(k|k)} & 0 & \cdots & 0 \\ \frac{\partial \hat{y}(k+2|k)}{\partial \Delta \boldsymbol{d}(k|k)} & \frac{\partial \hat{y}(k+2|k)}{\partial \Delta \boldsymbol{d}(k+1|k)} & \cdots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ \frac{\partial \hat{y}(k+N_p|k)}{\partial \Delta \boldsymbol{d}(k|k)} & \frac{\partial \hat{y}(k+N_p|k)}{\partial \Delta \boldsymbol{d}(k+1|k)} & \cdots & \frac{\partial \hat{y}(k+N_p|k)}{\partial \Delta \boldsymbol{d}(k+N_p-1|k)} \end{bmatrix}. \tag{12}$$

Each element of $\boldsymbol{G}_d$ is a row vector when several measured disturbances are considered.

Rather than deriving analytical model derivatives, PNMPC evaluates these matrices numerically [4,9]. For example, the jth column of $\boldsymbol{G}$ can be approximated by

$$\boldsymbol{G}_{(:,j)} \simeq \frac{\hat{\boldsymbol{Y}}\left(\boldsymbol{\Delta u} + \varepsilon_u \boldsymbol{e}_j\right) - \boldsymbol{F}}{\varepsilon_u}, \tag{13}$$

where ε_u is a small input perturbation and $\boldsymbol{e}_j$ selects the jth future control increment. An equivalent procedure is applied to the measured disturbances,

$$\boldsymbol{G}_{d(:,j)} \simeq \frac{\hat{\boldsymbol{Y}}\left(\boldsymbol{\Delta d} + \varepsilon_d \boldsymbol{e}_j\right) - \boldsymbol{F}}{\varepsilon_d}. \tag{14}$$

The PNMPC prediction can, therefore, be written compactly as

$$\hat{\boldsymbol{Y}} = \boldsymbol{F} + \boldsymbol{G}\boldsymbol{\Delta u} + \boldsymbol{G}_d \boldsymbol{\Delta \hat{d}}, \tag{15}$$

where $\boldsymbol{\Delta \hat{d}}$ denotes the assumed evolution of the measured disturbances along the prediction horizon. In this work, a persistence assumption is adopted: the most recent disturbance measurement is held constant over the complete prediction horizon. Thus,

$$\hat{\boldsymbol{d}}(k+j|k) = \boldsymbol{d}(k), \qquad j = 0, \ldots, N_p - 1. \tag{16}$$

Accordingly, only the disturbance increment available at the current sampling instant is considered,

$$\Delta\hat{\boldsymbol{d}}(k|k) = \boldsymbol{d}(k) - \boldsymbol{d}(k-1), \qquad \Delta\hat{\boldsymbol{d}}(k+j|k) = \mathbf{0}, \quad j = 1, \dots, N_p - 1. \tag{17}$$

The conventional PNMPC considered in this study incorporates the filtered prediction-error correction previously proposed for SCF applications in [11]. Defining the instantaneous model error as

$$e_m(k) = y(k) - \hat{y}(k), \tag{18}$$

the filtered correction is obtained as

$$e_f(k) = F_r(z^{-1})e_m(k), \tag{19}$$

where $F_r(z^{-1})$ is the selected prediction-error filter expressed in the discrete-time z-transform domain, where z denotes the complex transform variable and z^{-1} is used as the polynomial variable in the filter representation. This notation commonly represents discrete-time dynamic filters. For the CIESOL implementation considered in this work, the same filter adopted in [4] is used,

$$F_r(z^{-1}) = \frac{0.09516}{z - 0.9048}. \tag{20}$$

The corrected PNMPC output prediction is consequently

$$\tilde{\mathbf{Y}}_{\text{PNMPC}} = \hat{\mathbf{Y}} + \mathbf{1}_{N_p} e_f(k). \tag{21}$$

3.2. Learning-Based Practical Nonlinear Model Predictive Control

The LBPNMPC retains the nonlinear free and forced responses of the PNMPC formulation but replaces the fixed filtered-error correction in Equation (21) with the adaptive Oracle introduced in [4]. The Oracle estimates the future model error using a vector of plant pseudostates,

$$\begin{aligned} \boldsymbol{s}(k) = \big[&y(k-1), \dots, y(k-n_{y_s}), u(k-1), \dots, u(k-n_{u_s}), \\ &\boldsymbol{d}^\top(k-1), \dots, \boldsymbol{d}^\top(k-n_{d_s}), e(k-1), \dots, e(k-n_{e_s})\big]^\top \end{aligned} \tag{22}$$

where $e(k) = y(k) - \hat{y}(k)$ and n_{y_s}, n_{u_s}, n_{d_s}, and n_{e_s} define the corresponding regressor orders.

For each future instant $j = 1, \dots, N$, a different parameter vector $\boldsymbol{\Phi}_j$ is identified, yielding

$$O(k+j|k) = \hat{e}(k+j|k) = \boldsymbol{\Phi}_j^\top(k)\boldsymbol{s}(k). \tag{23}$$

The complete Oracle prediction is then arranged as

$$\boldsymbol{\Theta}(O(k)) = \boldsymbol{\Phi}^\top(k)\boldsymbol{s}(k) = \begin{bmatrix} O(k+1|k) \\ O(k+2|k) \\ \vdots \\ O(k+N|k) \end{bmatrix}. \tag{24}$$

The parameters in $\boldsymbol{\Phi}$ are updated online through the ER+RLS algorithm described in [4,24]. Separate parameter vectors are estimated for the different prediction instants, avoiding recursive propagation of a single one-step-ahead error model. When $N < N_p$, the last Oracle estimate is repeated from $N+1$ to N_p, following the formulation proposed in [4].

The LBPNMPC prediction is, therefore,

$$\tilde{\mathbf{Y}}_{\text{LBPNMPC}} = \underbrace{\boldsymbol{F} + \boldsymbol{G}\Delta\boldsymbol{u} + \boldsymbol{G}_d\Delta\hat{\boldsymbol{d}}}_{\hat{\mathbf{Y}}} + \boldsymbol{\Theta}(O(k)). \tag{25}$$

Hence, both controllers use the same nonlinear-model prediction and the same measured disturbances. Their distinction is restricted to the treatment of the residual prediction error: PNMPC applies the filtered correction $e_f(k)$ uniformly along the prediction horizon, whereas LBPNMPC uses an adaptive multistep error estimate obtained from recent plant information.

3.3. Optimisation Problem

Both controllers solve the same constrained quadratic optimisation problem. Using $\tilde{\boldsymbol{Y}}$ to denote either $\tilde{\boldsymbol{Y}}_{\text{PNMPC}}$ or $\tilde{\boldsymbol{Y}}_{\text{LBPNMPC}}$, the optimal future input increments are obtained from

$$\min_{\boldsymbol{\Delta u}} \quad J = \sum_{i=1}^{N_p} |\tilde{y}(k+i|k) - y_{\text{ref}}(k+i|k)|^2_{\boldsymbol{Q}} + \sum_{l=0}^{N_c-1} |\Delta u(k+l|k)|^2_{\boldsymbol{R}}, \tag{26}$$

subject to

$$\begin{aligned} u_{\min} \le u(k+l|k) \le u_{\max}, \quad & l = 0, \dots, N_c - 1, \\ \Delta u_{\min} \le \Delta u(k+l|k) \le \Delta u_{\max}, \quad & l = 0, \dots, N_c - 1, \\ y_{\min} \le \tilde{y}(k+i|k) \le y_{\max}, \quad & i = 1, \dots, N_p. \end{aligned} \tag{27}$$

Matrices $\boldsymbol{Q}$ and $\boldsymbol{R}$ weight the outlet-temperature tracking error and the variation of the manipulated HTF flow, respectively. The bounds on u, Δu, and y represent the operating and actuator constraints of the SCF.

Solving the constrained optimisation problem yields the optimal control-increment sequence $\Delta \boldsymbol{u}^*$. According to the receding-horizon strategy, only its first element, $\Delta u^*(k|k)$, is implemented at the current sampling instant. Hence, the manipulated input is updated as

$$u(k) = u(k-1) + \Delta u^*(k|k), \tag{28}$$

after which the measurements are updated, and the prediction and optimisation procedure is repeated.

3.4. Treatment of the Inlet-Temperature Disturbance

For the CIESOL SCF, the measured-disturbance vector contains solar irradiance, ambient temperature, and HTF inlet temperature,

$$\boldsymbol{d}(k) = \begin{bmatrix} I(k) & T_a(k) & T_{\text{in}}(k) \end{bmatrix}^\top. \tag{29}$$

Consequently, the effect of the measured inlet-temperature transient is explicitly included in the nominal prediction through $\boldsymbol{G}_d \boldsymbol{\Delta \hat{d}}$ for both controllers. As defined in Equation (16), the measured disturbances are assumed constant over the prediction horizon. Therefore, the controller has no prior knowledge of future variations in T_{in}; it propagates only the effect of the disturbance increment measured at the current sampling instant. At the following control update, the disturbance prediction is reconstructed using the new measurements.

The Oracle does not estimate or forecast T_{in}. Instead, it estimates the residual model error remaining after accounting for the nonlinear SCF model and the measured-disturbance contribution.

This distinction is central to the comparison performed in this work. Since PNMPC and LBPNMPC employ the same nonlinear model, disturbance measurements, transport delays, constraints, and optimisation problem, differences in disturbance-rejection performance arise from their different prediction-error compensation mechanisms.

4. Simulation Assessment

The simulation compares LBPNMPC with two PNMPC configurations under a clear-day condition selected to isolate the imposed T_{in} transients from high-frequency irradiance variations.

4.1. Simulation Setup

The simulation uses a measured clear-day dataset from the CIESOL facility comprising 7375 samples with a sampling period of $T_s = 1$ s. The measured solar irradiance and ambient-temperature profiles are retained, while the inlet-temperature trajectory is modified by introducing the disturbances described in Section 4.2. The initial operating conditions are $T_{\text{out}}(0) = 67.44\,^{\circ}\text{C}$ and $q(0) = 12.00\ \text{m}^3/\text{h}$.

The outlet-temperature reference comprises seven successive operating levels,

$$r = \{68, 70, 73, 76, 80, 84, 86\}\,^{\circ}\text{C}, \tag{30}$$

covering different thermal operating conditions.

The SCF controllers use the nominal model parameters reported in Table 1, including constant transport delays of $t_{d,T} = 41$ s for the inlet temperature and $t_{d,q} = 38$ s for the HTF flow. These values represent the nominal average delays used internally by all control configurations. To assess controller performance under transport-delay mismatch, the simulated plant considers a flow-dependent HTF delay [4]. In particular, the actual flow delay used in the simulation, denoted by $\hat{t}_{d_q}(q(t))$, is approximated as a piecewise function of the instantaneous HTF flow rate:

$$\hat{t}_{d_q}(q(t)) = \begin{cases} t_{d_q} + 10, & \text{if } 3 \le q(t) < 6, \\ t_{d_q} + 5, & \text{if } 6 \le q(t) < 9, \\ t_{d_q}, & \text{if } 9 \le q(t) < 12, \\ t_{d_q} - 10, & \text{if } 12 \le q(t). \end{cases} \tag{31}$$

The piecewise law in Equation (31) is adopted unchanged from the delay-mismatch benchmark studied in [4]. Its 28–48 s range represents a prescribed simulation approximation around the nominal 38 s delay. Here, this benchmark is retained to assess rapid inlet-temperature disturbance rejection; a new identification of the flow–delay relationship or a sensitivity study over alternative delay settings is outside the present scope.

Thus, while the controllers predict the plant behaviour using the nominal delay $t_{d,q} = 38$ s, the simulated plant is subject to a time-varying delay determined by the current flow rate. This creates a controlled plant–model mismatch that reproduces the dependence of HTF transport time on operating conditions and allows evaluation of the predictive controllers' robustness under varying-delay conditions.

Moreover, for LBPNMPC and PNMPC, measured disturbances are updated at every sampling instant and held constant over the prediction horizon according to Equation (16); neither controller accesses the future inlet-temperature trajectory.

Three control configurations are considered:

1. LBPNMPC: the complete learning-based formulation, including the measured-disturbance forced response $\boldsymbol{G}_d \Delta \hat{\boldsymbol{d}}$ and the adaptive Oracle prediction-error correction;
2. PNMPC: the conventional PNMPC with measured-disturbance feedforward and the filtered prediction-error correction described in Section 3.1;
3. PNMPC-noFF: a PNMPC formulation without disturbance feedforward. In this case, $\boldsymbol{G}_d$ is not computed and the forced response associated with the measured disturbance increment is omitted. The controller retains the same nonlinear free response and filtered prediction-error correction as the conventional PNMPC.

The distinction among the three prediction structures can, therefore, be expressed as

$$\tilde{\boldsymbol{Y}}_{\mathrm{LBPNMPC}} = \boldsymbol{F} + \boldsymbol{G}\Delta\boldsymbol{u} + \boldsymbol{G}_d\Delta\hat{\boldsymbol{d}} + \boldsymbol{\Theta}(O(k)), \tag{32}$$

$$\tilde{\boldsymbol{Y}}_{\mathrm{PNMPC}} = \boldsymbol{F} + \boldsymbol{G}\Delta\boldsymbol{u} + \boldsymbol{G}_d\Delta\hat{\boldsymbol{d}} + \mathbf{1}_{N_p} e_f(k), \tag{33}$$

$$\tilde{\boldsymbol{Y}}_{\mathrm{PNMPC-noFF}} = \boldsymbol{F} + \boldsymbol{G}\Delta\boldsymbol{u} + \mathbf{1}_{N_p} e_f(k). \tag{34}$$

The LBPNMPC–PNMPC comparison evaluates the implemented prediction-error-compensation schemes, including the excitation used for Oracle estimation, whereas PNMPC-noFF assesses the contribution of measured-disturbance feedforward.

Table 2 summarises the controller settings.

Table 2. Controller configurations used in the simulation study.

Setting	LBPNMPC	PNMPC	PNMPC-noFF
N_p	24	24	24
N_c	6	6	6
$\boldsymbol{Q}$	$40I$	$40I$	$40I$
$\boldsymbol{R}$	$10I$	$10I$	$10I$
$q_{\min}$ (m^3/h)	3	3	3
$q_{\max}$ (m^3/h)	15	15	15
Δq (m^3/h)	± 5	± 5	± 5
T_{out} bounds (°C)	10–100	10–100	10–100
$t_{d,T}$ (s)	41	41	41
$t_{d,q}$ (s)	38	38	38
Disturbance feedforward	Yes	Yes	No
$\boldsymbol{G}_d$	Computed	Computed	Not computed
Error correction	Oracle	Filtered error	Filtered error
n_{y_s}	2	–	–
n_{u_s}	2	–	–
n_{d_s}	2	–	–
n_{e_s}	2	–	–
Oracle horizon N	10	–	–
$\boldsymbol{R}_\infty$	$\frac{0.08}{1-\lambda}\boldsymbol{I}_{n_s}$	–	–
Prediction-error filter	–	$\frac{0.09516}{z-0.9048}$	$\frac{0.09516}{z-0.9048}$

The Oracle and filtered-error tunings in Table 2 follow those adopted for CIESOL in [4], employing $\lambda = 0.99$. The forgetting factor λ controls the estimator's memory, while $\boldsymbol{R}_\infty$ determines the exponential-resetting contribution. These settings influence adaptation to changing residual errors and sensitivity to noise. As in the previous implementation [4], a zero-mean white-noise dither with variance 0.03 is applied to LBPNMPC during simulation to maintain sufficient information for ER + RLS estimation; the PNMPC configurations do not require this excitation.

4.2. Inlet-Temperature Disturbance Scenarios

The simulated inlet-temperature disturbances are constructed from the actual measured T_{in} trajectory in Figure 3. Both the source record and the simulation are sampled at $T_{\mathrm{src}} = T_{\mathrm{sim}} = 1$ s. Both records use one-based sample indices, with elapsed times defined separately as $t_{\mathrm{src}} = (k_{\mathrm{src}} - 1)T_{\mathrm{src}}$ and $t_{\mathrm{sim}} = (k_{\mathrm{sim}} - 1)T_{\mathrm{sim}}$.

Five measured T_{in} valleys are used as templates. For a source interval $[k_a, k_b]$, the local baseline is the straight line between the mean of the 20 samples immediately before k_a and the mean of the 20 samples immediately after k_b. After subtracting this baseline, the residual is smoothed with a five-sample moving mean. A straight line joining the first and last smoothed residuals is then subtracted, positive residuals are clipped to zero,

and the resulting valley is divided by its negative minimum. The first and last normalised samples are explicitly set to zero. Each processed template has zero endpoints and a unit positive peak before multiplication by the prescribed negative amplitude. The templates are applied sample by sample on the simulation grid.

Table 3 reports the exact source windows and processed measured peak decreases. The temporal span is the difference between the first and last sample times; hence an interval containing N inclusive samples spans $(N - 1)$ s.

Table 3. Provenance of the five measured inlet-temperature templates. The source-window column gives one-based inclusive sample indices followed by the corresponding t_{src} interval in seconds, where t_{src} is elapsed time from the first sample. The reported peak is obtained after baseline removal, smoothing, endpoint correction, and clipping, but before unit-peak normalisation.

Template	Source Window: Samples; t_{src} (s)	N	Span (s)	Peak Decrease (°C)
1	55–217; 54–216	163	162	−1.1202
2	2484–2631; 2483–2630	148	147	−1.1282
3	3274–3412; 3273–3411	139	138	−1.3158
4	6117–6221; 6116–6220	105	104	−1.4875
5	6726–6878; 6725–6877	153	152	−1.8269

These processed measured decreases determine only the normalising factors. For simulated event j, the selected unit-peak template is multiplied by the independently prescribed amplitude a_j; therefore, the imposed peak decrease is exactly a_j, not the measured decrease in Table 3. The source windows identify portions of the temperature record selected for their profile shapes, not independently measured valve-opening durations. Their use defines controlled inlet-temperature disturbance tests rather than a complete reconstruction of expansion-vessel hydraulics. The imposed inlet temperature acts as a thermal boundary condition in the SCF energy-balance model, with the HTF flow determined by the controller. Upstream mixing and the mass and energy balances of the expansion-vessel branch are not reconstructed.

Table 4 gives the implemented template assignment and schedule. Event windows include both zero-valued template endpoints. The gap for event $j > 1$ is defined as $t_{\mathrm{start},j} - t_{\mathrm{end},j-1}$ using the tabulated sample timestamps. A reference transition is classified as overlapping when its change sample lies within the inclusive event interval, including equality with either boundary.

Table 4. Implemented schedule of the fourteen simulated inlet-temperature disturbances. Start, end, span, and gap are in elapsed simulation seconds, with $t_{\mathrm{sim}} = 0$ at sample 1. "None" means that no reference change occurs inside or at a boundary of the inclusive disturbance window.

Event	Template	Peak (°C)	Start (s)	End (s)	Span (s)	Gap (s)	Reference-Transition Overlap
1	1	−1.0	299	461	162	–	None
2	2	−2.0	699	846	147	238	None
3	3	−3.0	1249	1387	138	403	None
4	4	−4.0	1699	1803	104	312	None
5	5	−5.0	2199	2351	152	396	None
6	1	−2.0	2749	2911	162	398	None
7	2	−4.0	3049	3196	147	138	None
8	3	−1.5	3649	3787	138	453	None
9	4	−3.0	4199	4303	104	412	None
10	5	−5.0	4649	4801	152	346	None
11	1	−2.5	5199	5361	162	398	None
12	2	−4.5	5649	5796	147	288	None
13	3	−3.5	6299	6437	138	503	None
14	4	−5.0	7099	7203	104	662	None

All inter-event gaps are positive (minimum 138 s), confirming that the implemented disturbance windows do not overlap. Reference changes occur at $t_{\mathrm{sim}} = 1000, 2000, 3000,$

4000, 5000, and 6000 s; all six occur in inter-event gaps, and none coincides with an event boundary. The disturbances, therefore, span the seven reference levels without overlapping a reference-temperature transition.

Figure 4 compares the complete closed-loop responses. The principal comparison is between the matched-feedforward LBPNMPC and PNMPC cases.

As seen in Figure 4, PNMPC-noFF recovers more slowly from the imposed events and accumulates larger tracking errors over the complete simulation. Although receding-horizon updates let it react to new measurements, output deviations persist longer without the explicit disturbance contribution. Its behaviour, therefore, demonstrates the consequence of omitting disturbance feedforward while retaining feedback correction.

Between the two matched-feedforward controllers, LBPNMPC's time-varying multi-step error correction reduces outlet-temperature overshoot relative to PNMPC's horizon-wide filtered correction, as shown in the zoomed region of Figure 4. The difference is most evident during rapid T_{in} recovery, when disturbance evolution interacts with delayed HTF-flow action. In this phase, the Oracle correction represents an evolving residual-error profile rather than applying one filtered value uniformly across the horizon, yielding an earlier return to the tracking band.

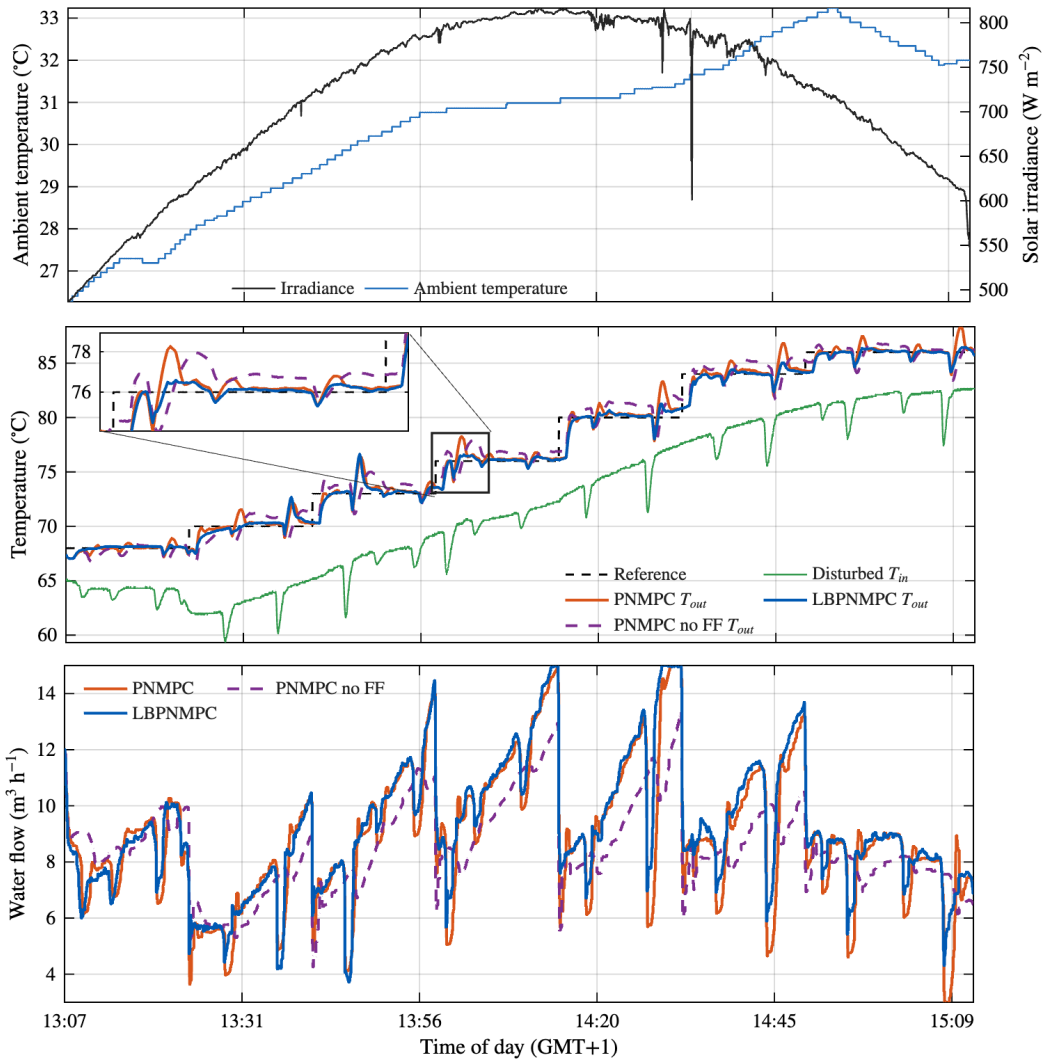


Figure 4. Closed-loop comparison under the clear-day inlet-temperature disturbance scenario. (**Top**) solar irradiance and ambient temperature. (**Middle**) outlet-temperature reference, simulated outlet temperatures, and disturbed HTF inlet temperature. (**Bottom**) manipulated HTF flow for LBPNMPC, PNMPC, and PNMPC-noFF.

4.3. Quantitative Performance Assessment

The performance indices defined in Section 2.3 are evaluated both over the complete simulation and over the response intervals associated with the imposed inlet-temperature disturbances. Table 5 summarises the results.

Table 5. Performance indices for the simulated inlet-temperature disturbance scenario. Lower values indicate better performance for all metrics. The mean recovery time index is computed in the 13 disturbance profiles considering the available time window.

Metric	LBPNMPC	PNMPC	PNMPC-noFF
Full-run IAE (°C s)	3511.43	4306.23	6367.30
Full-run ISE (°C^2 s)	5090.62	6670.99	8165.44
T_{in}-event IAE (°C s)	1025.49	1561.45	1457.87
T_{in}-event ISE (°C^2 s)	1258.95	2484.72	1725.87
Flow total variation (m^3/h)	212.20	239.55	107.22
Mean recovery time (s)	40.15	52.77	190.17

For event i, the recovery time is measured from $t_{\mathrm{end},i} + 41$ s, where $t_{\mathrm{end},i}$ is the end of the imposed inlet-temperature disturbance and the additional offset accounts for transport delay. Recovery occurs at the first subsequent instant when the outlet-temperature tracking error enters the ± 0.5 °C band and remains within it for at least 60 s. The observation interval ends immediately before the next disturbance onset shifted by the same delay, or at the end of the record. The qualifying 60-s interval must lie entirely within this observation interval. Calculations use the 1 s simulation sampling interval. Unrecovered events are excluded from the mean:

$$\bar{t}_{\mathrm{rec}} = \frac{1}{N_{\mathrm{rec}}} \sum_{i \in \mathcal{R}} t_{\mathrm{rec},i},$$

where $\mathcal{R}$ contains the recovered events and $N_{\mathrm{rec}} = |\mathcal{R}|$. The reference changes do not terminate the search and may influence the reported recovery times.

Compared with PNMPC, LBPNMPC reduces the full-run IAE and ISE by 18.5% and 23.7%, respectively. The improvement becomes more pronounced when the analysis is restricted to the inlet-temperature events: the event IAE is reduced by 34.3% and the event ISE by 49.3%. LBPNMPC also reduces the mean recovery time from 52.77 s to 40.15 s, a 23.9% improvement. These reductions are achieved alongside an 11.4% decrease in HTF-flow total variation, indicating that improved disturbance rejection does not require greater cumulative control movement.

The PNMPC-noFF results provide complementary information. Omitting the disturbance-forced response increases the full-run IAE and ISE to 6367.30 °C s and 8165.44 °C^2 s, respectively, while the mean recovery time increases to 190.17 s. Compared with this configuration, LBPNMPC reduces the full-run IAE by 44.9%, the full-run ISE by 37.7%, and the mean recovery time by 78.9%. The no-FF controller nevertheless produces the smallest flow total variation, indicating a less active manipulated-input trajectory.

The event IAE and ISE obtained with PNMPC-noFF are lower than those obtained with conventional PNMPC, although its full-run tracking performance and recovery time are substantially worse. The event windows follow the imposed inlet-temperature valleys and do not necessarily contain the complete outlet-temperature response. Delayed excursions and tracking errors persisting beyond these windows are excluded from the local integrals, so apparently favourable event indices can coexist with prolonged deviations and poorer complete-record performance. Therefore, these event integrals should not be interpreted independently as evidence of superior disturbance rejection. The result highlights the

importance of considering transient-error magnitude alongside recovery characteristics and the associated control effort.

Table 6 reports the event-resolved results for the two matched-feedforward controllers. Event IAE integrates the outlet-temperature tracking error over the prescribed inlet-temperature valley, from its onset through its return to baseline.

Table 6. Per-event performance of conventional filtered-error PNMPC and LBPNMPC. IAE: outlet-temperature event integral absolute error; TV: HTF-flow total variation; t_{rec}: recovery time. NR denotes no qualifying 60-s recovery interval within the observation window. IAE and TV are summed over all fourteen events, whereas recovery times are averaged over the thirteen recovered events for each controller; event 6 is unrecovered for both.

Event	PNMPC			LBPNMPC		
	IAE (°C s)	TV (m^3/h)	t_{rec} (s)	IAE (°C s)	TV (m^3/h)	t_{rec} (s)
1	28.71	3.73	0	21.06	3.73	0
2	67.67	4.19	0	29.36	3.75	0
3	97.31	4.05	14	37.20	3.09	0
4	110.94	4.67	40	132.01	5.01	80
5	179.01	9.97	139	186.05	10.07	61
6	67.91	7.23	NR	53.78	5.77	NR
7	174.14	6.27	64	75.16	5.08	42
8	59.25	4.14	3	39.64	3.20	0
9	71.15	4.03	7	37.48	3.06	68
10	193.40	16.35	263	104.13	14.53	269
11	71.95	7.34	19	62.26	5.06	0
12	215.45	7.74	86	104.62	5.84	2
13	102.74	6.40	8	63.84	4.14	0
14	121.79	6.46	43	78.88	2.23	0
Sum/mean	1561.45	92.59	52.77	1025.49	74.57	40.15

5. Experimental Validation at CIESOL

The simulation results are complemented by a CIESOL test of LBPNMPC under real operating conditions. Several inlet-temperature variations suggest the previously studied expansion-vessel fluid-exchange events, while the largest transient results from a separate operational intervention. This section assesses the controller response to these measured variations.

5.1. Experimental Implementation

The LBPNMPC was implemented in the existing CIESOL control architecture previously described in [4]. The predictive controller operates as an outer control loop that determines the HTF flow reference, while an inner PI controller with anti-windup regulates the circulation-pump frequency to achieve the commanded flow rate. The PI controller is implemented in the ideal form, $C_{\mathrm{PI}}(s) = K_p\left(1 + \frac{1}{T_i s}\right)$, with proportional gain $K_p = 0.571$, integral time $T_i = 1.775$ s, sampling period $T_{s,\mathrm{in}} = 1$ s and an effective delay of $T_d = 3$ s due to the flow-meter sensor lag. The desired-response parameter used for tuning was $\lambda_{\mathrm{PI}} = 0.887$ s. Anti-windup is implemented by correcting the integral accumulator using the previous sample's saturation error, with coefficient $T_r = 1.332$. The HTF-flow reference supplied to this loop is constrained to the range 2.5–11.7 m^3/h.

The experiment used a nominal acquisition period of 1 s, while the LBPNMPC control action was updated every 5 s. The controller used prediction and control horizons $N_p = 26$ and $N_c = 7$, respectively. The tracking- and control-move weighting matrices were

set to $\boldsymbol{Q} = 80\boldsymbol{I}$ and $\boldsymbol{R} = 10\boldsymbol{I}$. During on-site tuning, the tracking weight was adjusted relative to simulation to give greater priority to outlet-temperature regulation under the experimental conditions, which included a 5 s control update period, an inner PI flow loop with a 3 s effective sensor lag, and a tighter flow-increment limit of ±2.5 m^3/h per update. The experimental Oracle and ER+RLS settings were identical to those in Table 2: $n_{y_s} = n_{u_s} = n_{d_s} = n_{e_s} = 2$, $N = 10$, $\lambda = 0.99$, and $\boldsymbol{R}_\infty = 8\boldsymbol{I}_{n_s}$.

The HTF-flow command was constrained between 2.5 and 11.7 m^3/h, with maximum variations of ±2.5 m^3/h per control update, while the outlet-temperature prediction was constrained between 10 and 100 °C. These settings differ slightly from those used in the simulation because iterative trials were used to tune the controller to the real facility and its operating limitations. Following the experimental implementation in [4], no additional dither excitation was applied.

As in the simulation study, the measured disturbances were updated at every controller iteration and assumed constant over the prediction horizon. Consequently, the experimental LBPNMPC did not forecast future T_{in}, irradiance, or ambient-temperature variations.

The nonlinear model and Oracle formulation follow the LBPNMPC structure presented in Section 3.

5.2. Experimental Disturbance-Rejection Results

Figure 5 shows the complete experimental response as the outlet-temperature reference increases from approximately 51 to 58°C. Several smaller T_{in} transients are associated with operator-observed opening of the expansion-vessel connection valve. Their approximate intervals were identified retrospectively from the temperature record using the qualitative procedure described in Section 2.2. The dominant transient was produced by a separate operational intervention.

The operational intervention produces a pronounced inlet-temperature disturbance: T_{in} decreases from approximately 47.65 °C to 38.38 °C, a 9.27 °C drop that exceeds the 1–5 °C disturbances simulated. The corresponding measured outlet-temperature decrease is from 53.81 °C to 48.21 °C, i.e., 5.60 °C or approximately 60% of the inlet-temperature variation, while LBPNMPC adjusts the HTF flow. Because no matched experimental benchmark is available, this ratio is reported as an observation rather than a controller-specific attenuation measure.

The closed-loop response remains bounded, and T_{out} subsequently returns to the reference. It reaches and remains within the prescribed ±0.5 °C band 480 s after disturbance onset, demonstrating practical recovery from a severe inlet-temperature transient.

Finally, the commanded and measured flow trajectories in Figure 5 show small transient differences, consistent with the flow-measurement delay. The residual flow-tracking errors and inner-loop dynamics contribute to plant—model mismatch and can affect outlet-temperature prediction during rapid transients. These effects form part of the residual prediction error addressed by the Oracle. The inner-loop bandwidth and quantitative flow-tracking error were not separately evaluated in this study. Consequently, neglecting the inner-loop dynamics in the prediction model is treated as a modelling approximation, whose effect on disturbance-rejection performance has not been separately quantified.

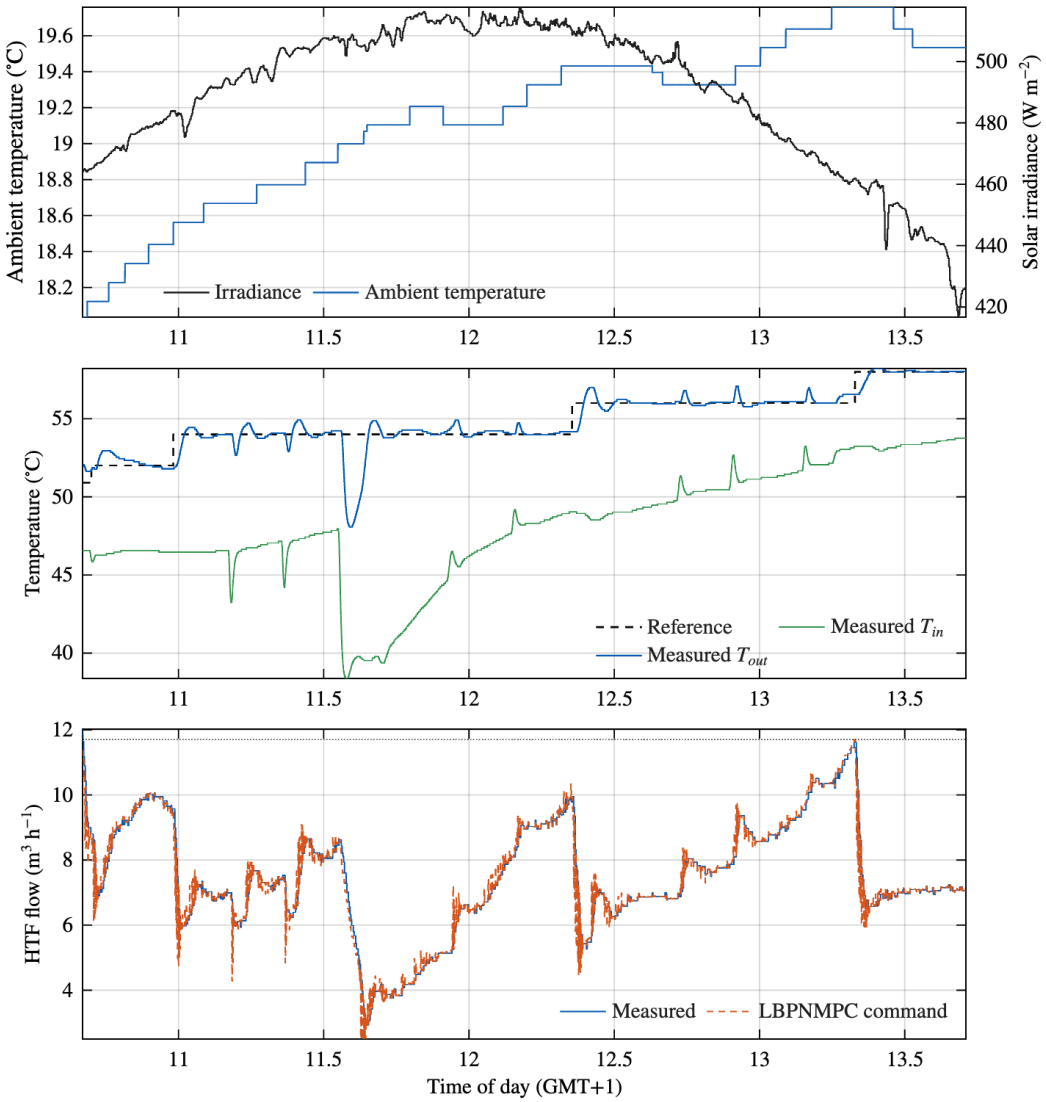


Figure 5. Experimental implementation of LBPNMPC in the CIESOL SCF. (**Top**) solar irradiance and ambient temperature. (**Middle**) outlet-temperature reference, measured SCF outlet temperature, and inlet temperature. (**Bottom**) commanded and measured HTF flow rate. The smaller inlet-temperature event intervals are approximate, with unquantified timing uncertainty (see Section 2.2).

6. Discussion

This section interprets the simulation and experimental results, emphasising the disturbance-rejection mechanism, the role of adaptive prediction-error correction, and the practical implications of the proposed approach.

6.1. Disturbance-Rejection Mechanism

The results identify rapid T_{in} recovery, rather than the initial decrease alone, as the most demanding phase: delayed flow compensation may persist as T_{in} recovers and thereby promote an outlet-temperature excursion, consistent with the physical mechanism established above. This timing explains why disturbance-event metrics and recovery time reveal differences that are less apparent from reference tracking over the complete record.

LBPNMPC and PNMPC use the same nonlinear model, measured-disturbance feedforward, delays, constraints, and optimisation problem, and neither accesses future T_{in}. Their different responses, therefore, reflect residual prediction-error correction: predefined filter dynamics in PNMPC and an Oracle adapted to recent plant evolution in LBPNMPC.

During fast T_{in} recovery, the smaller event errors and faster recovery obtained with LBPNMPC coincide with lower HTF-flow total variation. Thus, the improved matched-simulation response is not attributable to more cumulative control movement. The metric combination also indicates that the adaptive correction reduces the transient excursion without merely shifting the error into a longer settling interval.

Figure 6 compares the model prediction error with the Oracle estimate. The estimate follows the main residual-error variations around inlet-temperature recovery. While $\boldsymbol{G}_d\boldsymbol{\Delta}\hat{\boldsymbol{d}}$ represents the measured-disturbance effect predicted by the nonlinear model, the Oracle compensates for the remaining model–process discrepancy. This agreement supports the proposed interpretation, although it does not identify which unmodelled physical effect produces each residual component.

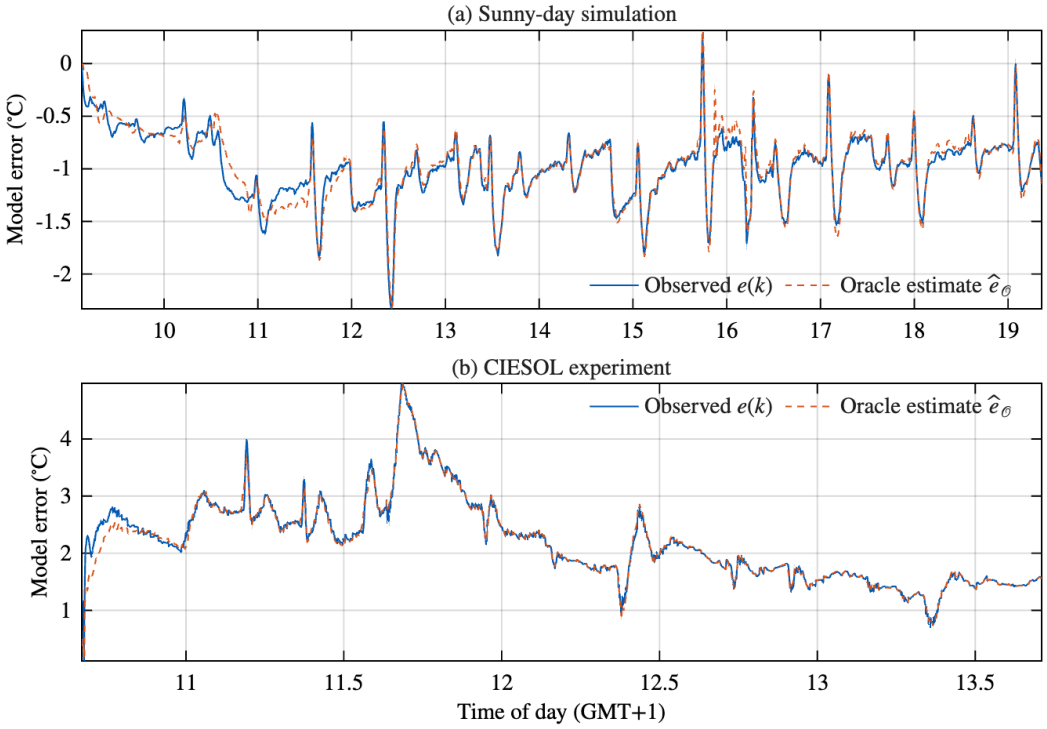


Figure 6. Measured model prediction error at time k, $e(k)$, compared with the one-step-ahead Oracle estimate $\hat{e}(k|k-1)$ computed at the previous control instant: (**a**) simulation and (**b**) CIESOL experiment. The comparison evaluates the Oracle's ability to predict the residual model error before the measurement at time k becomes available.

The Oracle is, therefore, neither an estimator nor a filter of T_{in}, which is measured directly and incorporated into the nonlinear model. It adaptively corrects remaining prediction inaccuracies arising from effects such as nonlinear hydraulics, mixing, and transport dynamics not fully captured by the nominal model.

PNMPC-noFF further demonstrates the relevance of disturbance feedforward. Its smallest flow total variation accompanies substantially larger full-run errors and much longer recovery, so the reduced activity indicates a weaker response to disturbance evolution rather than better rejection. Its event-integral values also show why local error indices should be interpreted together with recovery and complete-record performance.

6.2. Simulation and Experimental Interpretation

The studies provide complementary evidence. Matched simulations establish the quantitative comparison among LBPNMPC, PNMPC, and PNMPC-noFF. The CIESOL experiment instead assesses LBPNMPC during actual operation, including smaller varia-

tions from operator-observed expansion-vessel fluid exchange and a larger transient from a separate operational intervention.

An important mechanism contributing to this behaviour is the flow-dependent HTF transport delay discussed in [26] and described in the previous sections. As the manipulated flow $q(t)$ changes, the fluid residence time through the collector field also changes, modifying the delay between the applied control action and its effect on T_{out}. This effect is particularly relevant during inlet-temperature transients: when T_{in} decreases, the controller reduces $q(t)$ to compensate for the lower thermal level entering the field; however, the lower flow simultaneously increases the HTF transport delay. Consequently, when T_{in} rapidly recovers, the previous flow reduction may still propagate through the SCF, favouring the subsequent outlet-temperature excursion. The simulation scenarios explicitly represent this varying-delay condition, while the controllers retain the nominal average delay in their prediction models. The resulting plant–model mismatch, therefore, constitutes a main source of prediction error during the transient and helps explain why adaptive error compensation in LBPNMPC improves the closed-loop response relative to fixed filtered-error correction. The comparison, therefore, evaluates adaptive versus filtered prediction-error correction under the delay-mismatch benchmark inherited from [4]; it does not separately quantify the contribution of delay variability to the reported improvements.

For the separate operational intervention, the 9.27 °C inlet-temperature decrease corresponds to a measured 5.60 °C outlet-temperature decrease, approximately 60% of the inlet variation, followed by recovery towards the reference. This observation demonstrates bounded practical operation but does not isolate LBPNMPC's causal contribution to the attenuation.

A direct numerical comparison between simulation and experiment is inappropriate. The real controller was fine-tuned experimentally and uses slightly different horizons, weights, and constraints; the disturbance is larger, meteorological conditions differ, and the event occurs near a reference transition. Moreover, no matched experimental PNMPC benchmark is available. The experiment, therefore, demonstrates practical controller operation, not comparative superiority or direct reproduction of the simulation.

6.3. Practical Implications and Limitations

The results indicate complementary roles for measured-disturbance feedforward and adaptive residual-error correction. The former propagates the current measurement's modelled effect over the prediction horizon; the latter addresses transient dynamics not fully represented by the nominal model. Their combination provides the best simulated balance among rejection, recovery, and control activity. This distinction is important because improving the disturbance model and improving residual-error prediction address different sources of prediction inaccuracy. Since the persistence assumption uses only current measurements, the comparison concerns measured-disturbance response and not an advantage from forecasting future T_{in}.

The present results should nevertheless be interpreted within the study's scope. The matched-controller comparison uses one clear-day meteorological profile. This condition was deliberately selected to isolate the effect of the inlet-temperature disturbance, but the interaction between rapid T_{in} variations and high-frequency irradiance fluctuations remains to be investigated. Furthermore, the fourteen simulated disturbances are generated from five experimentally observed profiles and should, therefore, be regarded as controlled variations of representative disturbance shapes rather than independent experimental events. The percentage improvements are specific to the tested benchmark; their persistence under cloudy conditions, different initial states, or disturbance shapes outside the five templates has not been established. Uncertainty in event timing and the absence

of dedicated hydraulic measurements limit causal attribution of the observed transients to expansion-vessel exchange. Pump switching, other valve manoeuvres, or concurrent irradiance variations cannot be systematically excluded.

Finally, although the LBPNMPC formulation is unchanged from [4], its tuning still depends on knowledge of the process dynamics and desired closed-loop behaviour. In the experimental implementation, some control parameters were empirically fine-tuned to adapt the controller to the real facility. Developing more systematic procedures for selecting the MPC and Oracle parameters would, therefore, further facilitate application of the strategy to other solar thermal plants. Moreover, the reported comparative improvements are conditional on the adopted delay law, MPC tuning, and excitation used for Oracle estimation; their sensitivity to alternative delay settings, controller tunings, and excitation amplitudes has not been separately quantified.

7. Conclusions

This work evaluated the previously developed LBPNMPC for rejecting rapid HTF inlet-temperature disturbances in solar collector fields. A representative disturbance profile associated with expansion-vessel fluid exchange was first characterised using CIESOL operating data and then reproduced in simulation through measured-shape inlet-temperature valleys. In the matched clear-day simulation, LBPNMPC outperformed conventional filtered-error PNMPC, reducing the inlet-event IAE and ISE by 34.3% and 49.3%, respectively, while decreasing the mean recovery time by 23.9% and the HTF-flow total variation by 11.4%. Under the tested simulation conditions, these results demonstrate improved disturbance rejection with lower cumulative HTF-flow variation through adaptive Oracle correction.

The additional PNMPC-noFF case showed the importance of explicitly accounting for measured disturbances through the forced-response matrix $\boldsymbol{G}_d$. Although this configuration produced a less active flow variation, it resulted in considerably larger full-run tracking errors and a much slower recovery. The comparison, therefore, indicates that effective rejection of rapid T_{in} disturbances benefits from combining disturbance feedforward with adaptive prediction-error compensation.

Experimental operation at CIESOL further demonstrated the LBPNMPC response to a separate operational intervention producing a 9.27 °C decrease in T_{in}, substantially larger than the simulated disturbances. The corresponding measured T_{out} decrease was 5.60 °C, approximately 60% of the inlet-temperature variation, followed by recovery towards the reference. The observed Oracle behaviour was consistent with adaptive residual-error prediction during the transient.

The experimental results demonstrate the practical applicability of LBPNMPC under a severe inlet-temperature disturbance, while the matched simulation provides the quantitative comparison with conventional PNMPC. Future work should include repeated experiments comparing LBPNMPC with alternative predictive and machine-learning-based control strategies under matched operating and disturbance conditions, tests with combined inlet-temperature and irradiance disturbances, and more systematic procedures for tuning the MPC and Oracle parameters.

Author Contributions: I.M.L.P.: Conceptualisation, Formal analysis, Software, Investigation, Methodology, Writing—original draft, Writing—review and editing. J.D.G.: Conceptualisation, Formal analysis, Funding acquisition, Software, Investigation, Methodology, Writing—original draft, Writing—review and editing. J.L.G.: Conceptualisation, Formal analysis, Funding acquisition, Software, Investigation, Methodology, Writing—original draft, Writing—review and editing. M.B.: Conceptualisation, Formal analysis, Investigation, Methodology, Writing—review and editing. T.A.: Conceptualisation, Formal analysis, Investigation, Methodology, Funding acquisition, Writing—review and editing. All authors have read and agreed to the published version of the manuscript.

Funding: This publication is part of grant JDC2024-053231-I, funded by MICIU/AEI/10.13039/50110 011033 and by the ESF+. Publication of this paper was funded by the Spanish Ministry of Science, within the following projects: PID2023-150739OB-I00 and PDC2025-165379-I00 and DGP_PIDI_2024_0221 supported by the Junta de Andalucía (Consejería de Universidad, Investigación e Innovación) co-funded by the European Regional Development Fund (ERDF/FEDER) under the Andalusia FEDER Programme 2021–2027. The authors acknowledge support from grant PID2025-173578OB-I00 funded by MICIU/AEI/10.13039/501100011033.

Data Availability Statement: The simulation and experimental datasets and analysis scripts supporting this study are available from the corresponding author upon reasonable request.

Acknowledgments: During preparation of this manuscript, the authors used OpenAI Codex (CLI version 0.167.0-alpha.4) to assist with result plotting and coding consistency checks. The authors reviewed and edited the output and take full responsibility for the manuscript content. The authors declare no commercial interests related to the use of OpenAI Codex in this work.

Conflicts of Interest: The authors declare no conflicts of interest.

Abbreviations

The following abbreviations are used in this manuscript:

ER + RLS	Recursive least squares with exponential resetting
HTF	Heat-transfer fluid
IAE	Integral absolute error
ISE	Integral squared error
LBPNMPC	Learning-based practical nonlinear model predictive control
MPC	Model predictive control
PNMPC	Practical nonlinear model predictive control
SCF	Solar collector field
TV	Total variation

References

1. Johansen, T.A.; Storaa, C. Energy-based control of a distributed solar collector field. *Automatica* **2002**, *38*, 1191–1199. [CrossRef]
2. Johansen, T.A.; Hunt, K.J.; Petersen, I. Gain-scheduled control of a solar power plant. *Control Eng. Pract.* **2000**, *8*, 1011–1022. [CrossRef]
3. Barão, M.; Lemos, J.M.; Silva, R.N. Reduced complexity adaptive nonlinear control of a distributed collector solar field. *J. Process Control* **2002**, *12*, 131–141. [CrossRef]
4. Pataro, I.M.L.; Gil, J.D.; Álvarez, J.D.; Guzmán, J.L.; Lemos, J.M.; Berenguel, M. Learning-based practical nonlinear predictive controller for solar thermal collector fields. *IEEE Trans. Control Syst. Technol.* **2025**, *33*, 2108–2120. [CrossRef]
5. Gil, J.D.; Roca, L.; Zaragoza, G.; Normey-Rico, J.E.; Berenguel, M. Hierarchical control for the start-up procedure of solar thermal fields with direct storage. *Control Eng. Pract.* **2020**, *95*, 104254. [CrossRef]
6. Camacho, E.F.; Bordons, C. *Model Predictive Control*; Springer: London, UK, 2007.
7. Morari, M.; Lee, J.H. Model predictive control: Past, present and future. *Comput. Chem. Eng.* **1999**, *23*, 667–682. [CrossRef]
8. Qin, S.J.; Badgwell, T.A. A survey of industrial model predictive control technology. *Control Eng. Pract.* **2003**, *11*, 733–764. [CrossRef]
9. Plucenio, A.; Pagano, D.; Bruciapaglia, A.; Normey-Rico, J.E. A practical approach to predictive control for nonlinear processes. *IFAC Proc. Vol.* **2007**, *40*, 210–215. [CrossRef]
10. Andrade, G.A.; Pagano, D.J.; Álvarez, J.D.; Berenguel, M. A practical NMPC with robustness of stability applied to distributed solar power plants. *Sol. Energy* **2013**, *92*, 106–122. [CrossRef]
11. Pataro, I.M.L.; Gil, J.D.; Guzmán, J.L.; Berenguel, M.; Cañadas, I. Predictive control strategies for solar furnace systems on the basis of practical constrained solutions. *J. Process Control* **2023**, *132*, 103114. [CrossRef]
12. Gao, X.h.; Wei, S.; Su, Z.g. Optimal model predictive rejection control for nonlinear parabolic trough collector with lumped disturbances. *Trans. Inst. Meas. Control* **2021**, *43*, 1903–1914. [CrossRef]
13. Pannocchia, G. Offset-free tracking MPC: A tutorial review and comparison of different formulations. In Proceedings of the 2015 European Control Conference (ECC), Linz, Austria, 15–17 July 2015; pp. 527–532. [CrossRef]
14. Pannocchia, G.; Rawlings, J.B. Disturbance models for offset-free model-predictive control. *AIChE J.* **2003**, *49*, 426–437. [CrossRef]

15. Chen, W.H.; Yang, J.; Guo, L.; Li, S. Disturbance-Observer-Based Control and Related Methods—An Overview. *IEEE Trans. Ind. Electron.* **2016**, *63*, 1083–1095. [CrossRef]
16. Alsharkawi, A.; Rossiter, J.A. Towards an Improved Gain Scheduling Predictive Control Strategy for a Solar Thermal Power Plant. *IET Control Theory Appl.* **2017**, *11*, 1938–1947. [CrossRef]
17. Lu, X.; Liang, Z. Application of Model Predictive Control Based on Kalman Filter in Solar Collector Field of Solar Thermal Power Generation. *Energy Eng.* **2021**, *118*, 1171–1183. [CrossRef]
18. Song, Y.; Xia, C.; Zhang, J.; Wang, J.; Li, Y. Disturbance rejection model predictive control for solar collector fields based on fast Fourier transform. In Proceedings of the 12th IFAC Symposium on Control of Power and Energy Systems—CPES 2024, Rabat, Morocco, 10–12 July 2024; Volume 58, pp. 284–289. [CrossRef]
19. Wei, S.; Gao, X.; Li, Y. Constraint Optimal Model-Based Disturbance Predictive and Rejection Control Method of a Parabolic Trough Solar Field. *Energies* **2024**, *17*, 5804. [CrossRef]
20. Zhang, X.; Wang, D.; Chen, C.; Song, X.; Zhao, S. Research on Stochastic Dual Model Predictive Control and Application to Solar Thermal Collector Field. *Electronics* **2025**, *14*, 2048. [CrossRef]
21. Carreño-Zagarra, J.; Poveda-Rodriguez, D.; Flórez, M. Robust Control of Irrigation Systems Using Predictive Methods and Disturbance Rejection. *Inventions* **2025**, *10*, 11. [CrossRef]
22. Hewing, L.; Wabersich, K.P.; Menner, M.; Zeilinger, M.N. Learning-Based Model Predictive Control: Toward Safe Learning in Control. *Annu. Rev. Control Robot. Auton. Syst.* **2020**, *3*, 269–296. [CrossRef]
23. Aswani, A.; Gonzalez, H.; Sastry, S.S.; Tomlin, C. Provably safe and robust learning-based model predictive control. *Automatica* **2013**, *49*, 1216–1226. [CrossRef]
24. Lai, B.; Bernstein, D.S. Exponential resetting and cyclic resetting recursive least squares. *IEEE Control Syst. Lett.* **2023**, *7*, 985–990. [CrossRef]
25. Lemos, J.M.; Neves-Silva, R.; Igreja, J. *Adaptive Control of Solar Energy Collector Systems*; Springer: London, UK, 2014. [CrossRef]
26. Normey-Rico, J.E.; Bordons, C.; Berenguel, M.; Camacho, E.F. A robust adaptive dead-time compensator with application to a solar collector field. In Proceedings of the IFAC-Workshop on Linear Time Delay Systems (LDTS'98), Grenoble, France, 6–7 July 1998; pp. 105–110. [CrossRef]

Disclaimer/Publisher's Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.