Research Square

Preprints are preliminary reports that have not undergone peer review.
They should not be considered conclusive, used to inform clinical practice,
or referenced by the media as validated information.

Cross-Coupling-Aware Model Predictive Control for Multi-Zone Thermal-Vacuum Chambers: Tracking Performance, Robustness, and Scalability

Xuan-Hai Le
Vietnam National University, Hanoi
Minh-Hoang Pham
phamminhhoang2004@gmail.com
Vietnam National University, Hanoi
Trong-Nam Do
Hanoi University of Science and Technology
Xuan-Minh Dinh
Phenikaa University

Research Article

Keywords:

Posted Date: November 24th, 2026

DOI: https://doi.org/10.21203/rs.3.rs-11645002/v1

License: This work is licensed under a Creative Commons Attribution 4.0 International License.
Read Full License

Additional Declarations: No competing interests reported.

Cross-Coupling-Aware Model Predictive Control for Multi-Zone Thermal-Vacuum Chambers: Tracking Performance, Robustness, and Scalability

Anonymous Author(s)

Affiliation withheld for double-anonymous review.

Abstract

Coordinated temperature regulation in multi-zone thermal-vacuum chambers (TVACs) is challenging because each heater influences several zones through dynamic thermal interactions. Comparisons of centralized and decentralized predictive controllers usually change the model, state information, and tuning simultaneously, so the contribution of interaction modeling itself remains unclear. This study examines when explicit interaction modeling benefits centralized model predictive control (MPC), using a physics-inspired synthetic benchmark: a heterogeneous nine-zone multiple-input multiple-output first-order-plus-dead-time (MIMO-FOPDT) model with 49 active heater-to-zone paths, spatially dependent dynamics, and adjustable coupling strength. A constrained coupling-aware centralized MPC (MIMO-MPC) is compared with decentralized proportional–integral (PI) control, decentralized MPC, and static-decoupled MIMO-PI in four simulation studies. With an exact nominal model and simulator state prediction, MIMO-MPC reaches a root-mean-square error (RMSE) of 0.0008 °C versus 0.0097 °C for decentralized MPC; this idealized gap combines cross zone model information with an information advantage. When both MPCs use the same measurement-based output correction and settings, they coincide without coupling, as predicted by an exact decomposition of the joint quadratic program, and the MIMO-MPC advantage grows with coupling strength to 12.8% at nominal coupling. Separately tuned MPCs differ by a factor of about six even without coupling, reflecting horizon selection rather than coupling awareness. The advantage persists under disturbances and ±10–20% parameter mismatch but shrinks markedly under measurement noise. Scalability tests on 9–36 zones preserve the controller ranking, while centralized optimization time grows faster. The results delimit when coupling-aware prediction is beneficial and provide a basis for experimental validation.

Keywords: Thermal-vacuum chamber, Model predictive control, Thermal coupling, Multi-zone temperature control, MIMO systems

1 Introduction

Thermal-vacuum (TVAC) testing is an important part of spacecraft development and qualification, providing a controlled environment in which hardware is exposed to prescribed thermal and vacuum conditions. The design of TVAC facilities and the effectiveness of thermal-vacuum tests have been examined for small satellites and microsatellites [1, 2]. During testing, heater commands must regulate temperatures across a test article while accommodating thermal inertia, delays, and heat exchange between regions. In multi-channel configurations, an individual heater may influence not only its intended zone but also neighboring zones. Consequently, independently regulating each heater–temperature pair may overlook interactions that affect coordinated tracking. Most TVAC temperature controllers reported to date are loop-wise designs whose treatment of inter-channel interaction is implicit. Proportional–integral–derivative (PID) tuning and optimization [3, 4], fuzzy and adaptive neural-network PID [5, 6], and neural-network identification with self-tuning PID [7] improve individual loops but do not predict how one heater affects other zones. For multi-channel systems, Li et al. grouped strongly interacting channels and combined neural-network PID with Smith prediction [8], and Chang et al. used average-temperature feedback to balance heating across channels [9]; in both cases interaction is mitigated by the loop structure rather than included in a multivariable prediction, so the effect of interaction modeling is not separated from that of the loop structure. Data-driven and hybrid models have also been proposed to predict temperatures in aerospace thermal tests [10], but they address open-loop prediction rather than closed-loop coordination. Other work has examined intelligent control systems and thermal regulation models for vacuum thermal testing [11, 12].

Predictive control provides an alternative framework because future temperature responses can be estimated from a dynamic model and actuator constraints can be included directly in the optimization [13–15]. Step-response-based formulations of this kind have been used in industrial multivariable control for decades [16]; the present study therefore does not propose a new Model Predictive Control (MPC) algorithm, but examines how much of its benefit on a coupled heater array is attributable to modeling the cross-channel dynamics. Park and Kim investigated nonlinear MPC for unit-level TVAC temperature regulation using a nonlinear heat-exchange model with radiative and conductive terms, whose parameters were obtained from actual TVAC test data; their study examined specimen temperature tracking and responses to internal heat generation [17]. In a recent preprint, Klupar et al. applied nonlinear MPC with an unscented Kalman filter to eight coupled heater nodes of a spacecraft optical payload and reported an RMSE reduction of about 88% relative to a single-sensor PI loop [18]. These studies demonstrate the relevance of predictive control in spacecraft thermal applications. However, because the predictive controllers differ from their baselines simultaneously in model structure, state estimation, actuator coordination,

and tuning, the share of the improvement that is due to modeling spatial coupling cannot be inferred from them.

The question of how much interaction information a multi-zone controller requires has been examined more systematically for buildings and other distributed thermal processes. Distributed MPC formulations have been developed for temperature regulation with coupled constraints [19], and subsequent studies have examined distributed and nonlinear predictive control for multi-zone buildings [20, 21]. Centralized, decentralized, and distributed MPC architectures differ in how much model and state information is shared among subsystems and in the associated computational and communication cost [22, 23]. Comparisons of centralized and decentralized MPC in building applications show that the controller architecture determines how shared model information is used to coordinate zones [24]. Zhao et al. incorporated inter-zone heat exchange into the predictor of an MPC for data-center cooling and evaluated it with and without the coupling terms [25], and a 2026 study of building districts showed that the coordination architecture, not only the controller type, determines tracking accuracy and the spatial distribution of control effort [26]. A recent review identifies multi-agent frameworks as an open direction for building MPC [27].

Three limitations of this body of work motivate the present study. First, architectures are usually compared with different models, estimators, and tuning, so the effect of interaction information is confounded with these factors. Second, coupling strength is rarely varied systematically; it is therefore unclear below which interaction level a decentralized design suffices. Third, the effects of incorporating dynamic cross-channel interaction information into the prediction model on control performance have not been systematically characterized for multi-zone TVAC systems, particularly in comparison with control architectures based on local prediction or static decoupling. Static decoupling relies only on the steady-state gain matrix and cannot capture heterogeneous cross-channel time constants and delays [28]. Building applications also differ from TVAC testing in their slow dynamics and economic objectives, so their conclusions do not transfer directly to the tight trajectory tracking required in thermal testing.

For a multi-input multi-output (MIMO) thermal plant, different control architectures retain different levels of interaction information. Decentralized PI and MPC may use only the direct heater-to-zone channels, whereas static decoupling uses the steady-state gain matrix to compensate for interactions without representing their distinct transient dynamics [28]. A centralized MIMO-MPC formulation can instead incorporate cross-channel gains, time constants, and delays into a joint prediction model and optimize all heater moves subject to amplitude and rate constraints. However, improved tracking in a nominal comparison does not, by itself, establish that dynamic coupling information is responsible: the state information used to initialize the prediction, reference preview, prediction and control horizons, tuning weights, and test conditions may also influence the observed differences.

The present study addresses these gaps with a controlled, reproducible evaluation on a TVAC-inspired synthetic benchmark. Specifically, this study evaluates the performance of coupling-aware centralized model predictive control by incorporating the dynamic characteristics of cross-zone thermal interactions into the prediction model.

Through comparisons with reference control architectures, the effects of thermal coupling strength, controller tuning, and operating conditions on temperature tracking accuracy, robustness, and computational cost are investigated. The benchmark is deliberately synthetic so that interaction strength and uncertainty can be varied independently; its results should therefore not be interpreted as experimental validation of a physical TVAC installation. The main contributions are as follows:

- *A parametric benchmark with controllable coupling.* A heterogeneous MIMO first-order-plus-dead-time (FOPDT) generator for rectangular heater arrays is specified, with explicit and physically motivated rules for nearest- and diagonal-neighbor gains, lags, and delays, and with a scalar factor that scales coupling strength without altering the direct-path dynamics. The nine-zone instance has 49 active heater-to-zone paths.
- *Separation of architectural and informational effects.* The state information available to each predictor is stated explicitly, and it is shown that, with a direct-path model, the joint MPC quadratic program decomposes exactly into independent local problems. Joint optimization alone therefore yields no benefit, and under matched information and settings any advantage of the centralized controller is attributable to the coupling terms of its model. This is verified by a matched-setting coupling sweep in which both architectures coincide at zero coupling and diverge progressively as coupling increases, whereas the large gaps observed under simulator-state prediction or separate tuning are shown to arise from information and horizon differences.
- *Conditions and costs of interaction modeling.* The advantage is quantified as a function of coupling strength, controller tuning, measurement noise, external load, parametric mismatch, actuator limits, and array size from 9 to 36 zones. The mechanism of the observed noise sensitivity is analyzed, and computational cost is reported together with the size of the optimization problem and percentile solution times.

The remainder of this paper is organized as follows. Section 2 describes the generalized thermal model and nine-zone benchmark together with the rationale and validity range of its parameters. Section 3 presents the reference controllers and the coupling-aware MIMO-MPC formulation. Section 4 defines the simulation protocols and performance measures. Section 5 discusses the tracking, robustness, constraint-handling, and scalability results, followed by the conclusions in Section 6.

2 Problem Statement and Thermal Plant Modeling

2.1 Thermal Dynamics of an Individual Heating Channel

In a thermal-vacuum chamber (TVAC), infrared heaters may be used to regulate the temperatures of multiple regions of a test article. The thermal response of each region depends on its effective thermal inertia, heat exchange with the surroundings, and heat spreading from adjacent regions. For an individual zone, a simplified energy balance can be written as

$$C_i \dot{T}_i(t) = Q_i(t) - Q_{\text{loss},i}(t), \tag{1}$$

where T_i is the zone temperature, C_i is an equivalent thermal capacitance, Q_i is the effective heating power, and $Q_{\text{loss},i}$ represents the net heat loss. This lumped balance provides physical motivation for the reduced-order model; it is not used as a separately calibrated nonlinear energy-balance model in the simulations.

The effective heater power may be related to a scaled command U_i through

$$Q_i(t) = \eta_i U_i(t), \tag{2}$$

where η_i is an effective conversion coefficient. Around a nominal operating condition, a direct heater-to-zone path is approximated by the first-order-plus-dead-time (FOPDT) transfer function

$$G_{ii}(s) = \frac{K_{ii}}{\tau_{ii}s + 1} \exp(-\theta_{ii}s), \tag{3}$$

where K_{ii}, τ_{ii}, and θ_{ii} are its steady-state gain, time constant, and effective delay, respectively. FOPDT-based descriptions have been considered in thermal-control and multi-channel TVAC studies [4, 5, 7, 8]. In the present work, the FOPDT coefficients are prescribed benchmark parameters rather than quantities identified from a particular TVAC installation.

2.2 Generalized MIMO Thermal Model with Spatial Interaction

Consider N heater inputs and N associated temperature zones arranged on a rectangular grid with n_r rows and n_c columns, such that $N = n_r n_c$. Let

$$\mathbf{U}(t) = [U_1(t), \ldots, U_N(t)]^\top, \qquad \mathbf{y}(t) = [T_1(t), \ldots, T_N(t)]^\top. \tag{4}$$

For a linearized description around nominal heater commands and temperatures, the input–output relation is

$$\delta\mathbf{y}(s) = \mathbf{G}_N(s)\delta\mathbf{U}(s), \qquad \mathbf{G}_N(s) \in \mathbb{R}(s)^{N\times N}. \tag{5}$$

Here, $\delta\mathbf{y}$ and $\delta\mathbf{U}$ denote deviations from the nominal operating condition, whereas $\Delta\mathbf{U}(k)$ is reserved for the sample-to-sample actuator increment introduced in Section 2.7. Each active path from heater i to zone j is represented as

$$G_{ji}(s) = \frac{K_{ji}}{\tau_{ji}s + 1} \exp(-\theta_{ji}s), \qquad (j,i) \in \mathcal{E}, \tag{6}$$

where $\mathcal{E}$ is the set of active heater-to-zone paths. For $(j,i) \notin \mathcal{E}$, $G_{ji}(s) = 0$. Rows of $\mathbf{G}_N$ correspond to temperature zones and columns to heater commands. The diagonal entries represent direct heating, whereas nonzero off-diagonal entries represent cross-zone interactions.

To define the spatial topology, assign each zone i a grid index (r_i, c_i), with $1 \le r_i \le n_r$ and $1 \le c_i \le n_c$. The sets of directed nearest-neighbor and diagonal-neighbor

paths are

$$\mathcal{E}_{\mathrm{N}} = \{(j,i) : |r_j - r_i| + |c_j - c_i| = 1\}, \tag{7}$$
$$\mathcal{E}_{\mathrm{D}} = \{(j,i) : |r_j - r_i| = 1,\ |c_j - c_i| = 1\}. \tag{8}$$

The complete interaction set is $\mathcal{E} = \{(i,i) : 1 \le i \le N\} \cup \mathcal{E}_{\mathrm{N}} \cup \mathcal{E}_{\mathrm{D}}$. Counting both directions of each cross-zone path gives

$$N_{\mathrm{near}} = 2\big[n_r(n_c - 1) + n_c(n_r - 1)\big], \tag{9}$$
$$N_{\mathrm{diag}} = 4(n_r - 1)(n_c - 1), \tag{10}$$
$$N_{\mathrm{active}} = N + N_{\mathrm{near}} + N_{\mathrm{diag}}. \tag{11}$$

These expressions define the channel topology for arbitrary rectangular grid dimensions; they do not require a fixed 9×9 transfer matrix.

The direct gains K_{ii} are assigned independently for each zone. For $i \neq j$, the steady-state cross-channel gains follow

$$K_{ji} = \begin{cases} \gamma_{\mathrm{N}}\sqrt{K_{ii}K_{jj}}, & (j,i) \in \mathcal{E}_{\mathrm{N}}, \\ \gamma_{\mathrm{D}}\sqrt{K_{ii}K_{jj}}, & (j,i) \in \mathcal{E}_{\mathrm{D}}, \\ 0, & \text{otherwise}, \end{cases} \tag{12}$$

where γ_{N} and γ_{D} are prescribed nearest-neighbor and diagonal-neighbor coupling coefficients. This construction produces symmetric steady-state cross-gains when the same coefficients are used in both directions, although the associated delays need not be symmetric.

For each active off-diagonal path, the benchmark construction assigns

$$\tau_{ji} = \kappa_\tau \max(\tau_{ii}, \tau_{jj}), \qquad \kappa_\tau = 1.2, \tag{13}$$
$$d_{ji} = d_{ii} + \delta_d, \qquad \delta_d = 1, \qquad \theta_{ji} = d_{ji}T_s, \tag{14}$$

where d_{ii} is the direct-path delay, in samples, of source heater i. Thus, cross-channel paths are slower than their corresponding direct paths and have one additional sample of delay relative to the source heater. These are modeling choices for the synthetic benchmark, not universal physical laws for TVAC systems; their rationale is given in Section 2.4.

2.3 Nine-Zone Benchmark Selection and Parameterization

The generalized formulation is instantiated as a 3×3 grid with nine heater inputs and nine controlled temperature zones. This is the smallest square grid that includes corner, edge, and interior locations, whose distinct numbers of neighbors provide heterogeneous local connectivity. Specifically, a corner zone has three neighboring zones, a non-corner edge zone has five, and the center zone has eight when horizontal, vertical, and diagonal adjacency are included. The 3×3 arrangement therefore offers a compact case for examining spatial coupling while keeping repeated controller

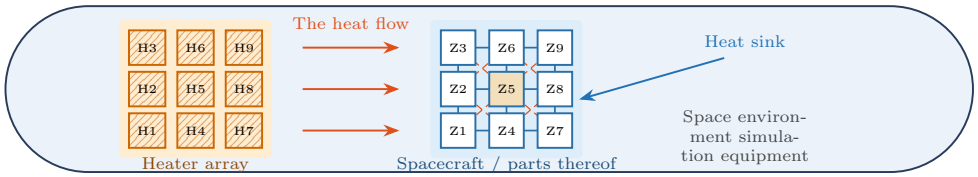

Fig. 1 Conceptual TVAC-inspired multi-zone thermal control configuration with a 3×3 heater array and nine interacting temperature zones. The schematic illustrates the application setting of the synthetic benchmark, not an experimentally identified chamber.

Table 1 Deterministic direct-channel parameters of the nine-zone synthetic benchmark.

Zone i	K_{ii}	τ_{ii} (s)	d_{ii} (samples)
1	0.91	68	2
2	1.08	74	3
3	0.97	83	2
4	1.16	91	4
5	1.25	105	3
6	0.88	78	5
7	1.11	116	2
8	1.03	87	3
9	0.94	98	4

tuning, uncertainty tests, and coupling sweeps computationally manageable. It is a selected benchmark, not a claim that nine zones constitute an optimal or universally representative TVAC configuration.

Fig. 1 shows the conceptual TVAC-inspired arrangement considered in the study. The model represents a heater array acting on a test article partitioned into nine thermal zones; the figure is illustrative rather than a depiction of an experimentally characterized apparatus.

For $n_r = n_c = 3$, Eqs. (9)–(11) give 24 directed nearest-neighbor paths, 16 directed diagonal-neighbor paths, and 9 direct paths. Hence, 49 of the 81 possible heater-to-zone channels are active. The nominal grid spans 1.20 m by 1.00 m; spatial coordinates establish the layout, whereas channel activation is determined by adjacency rather than by a Euclidean-distance cutoff.

The direct-path parameters are deterministic and are listed in Table 1. Their ranges are $K_{ii} \in [0.88, 1.25]$, $\tau_{ii} \in [68, 116]$ s, and $d_{ii} \in \{2, 3, 4, 5\}$. The nominal coupling coefficients are $\gamma_{\mathrm{N}} = 0.25$ and $\gamma_{\mathrm{D}} = 0.12$; Eq. (12) then determines all nonzero cross-channel gains. The time constants and delays of these paths follow Eqs. (13) and (14).

The resulting steady-state gain matrix is illustrated in Fig. 2. Its nonzero pattern indicates which paths are retained; the number of active paths alone does not quantify

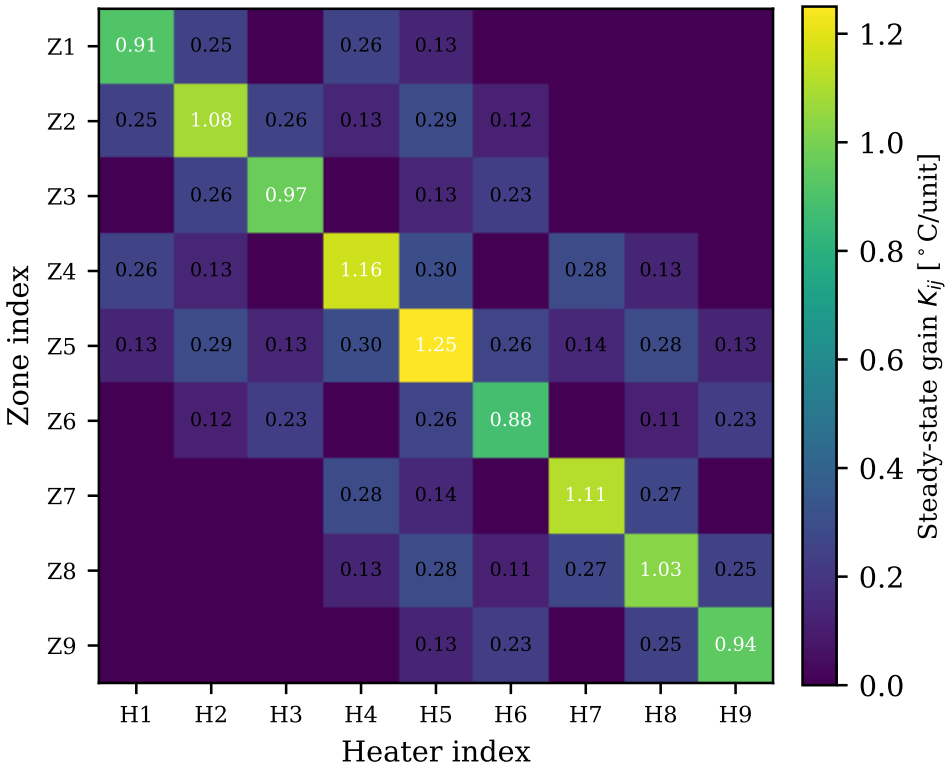

Fig. 2 Steady-state heater-to-zone gain matrix $\mathbf{K}$ for the nine-zone synthetic benchmark. Off-diagonal nonzero entries represent nearest-neighbor and diagonal-neighbor interactions.

the strength of thermal coupling. Coupling strength is therefore summarized by scale-free measures of the normalized gain matrix $\tilde{\mathbf{K}} = \mathbf{D}_K^{-1/2}\mathbf{K}\mathbf{D}_K^{-1/2}$, $\mathbf{D}_K = \mathrm{diag}(\mathbf{K})$: the off-diagonal Frobenius ratio $\|\tilde{\mathbf{K}} - \mathbf{I}\|_F / \|\mathbf{I}\|_F = 0.406$, the largest normalized off-diagonal gain 0.241, and the deviation of the relative gain array (RGA) from identity, 0.210. These values indicate moderate, but not negligible, interaction.

The nine-zone configuration is used for the detailed nominal, tuning, coupling-sensitivity, and robustness studies. Separate synthetic arrays with 16, 25, and 36 zones are examined in the scalability study to assess how controller performance and computational demand change with system dimension. Their parameter-generation rules and comparison settings are reported with the scalability results (Section 4.5), rather than presumed to be identical to those of the nine-zone plant.

2.4 Rationale for the Benchmark Parameters

The parameter rules are chosen to be qualitatively consistent with a linearized lumped thermal network, while remaining simple enough to be varied systematically.

Symmetric, geometric-mean cross-gains, Eq. (12). For a linear (or linearized) conductive–radiative network with heat inputs, $\mathbf{C}\dot{\mathbf{T}} = -\mathbf{L}\mathbf{T} + \mathbf{Q}$, the conductance matrix $\mathbf{L}$ is symmetric by reciprocity, so the steady-state map from heat inputs to temperatures, $\mathbf{L}^{-1}$, is also symmetric. Scaling the cross-gain by $\sqrt{K_{ii}K_{jj}}$ preserves this symmetry while making γ_{N} and γ_{D} dimensionless fractions of the local gain level, so that the same coefficients can be reused for arrays with different direct gains. The diagonal coefficient is smaller than the nearest-neighbor coefficient ($\gamma_{\mathrm{D}}/\gamma_{\mathrm{N}} = 0.48$) because diagonally adjacent zones share only a corner: their lateral conduction path is longer and their mutual radiative view factor is smaller. The magnitudes $\gamma_{\mathrm{N}} = 0.25$ and $\gamma_{\mathrm{D}} = 0.12$ were selected to yield the moderate interaction level reported above, for which the RGA remains close to identity and decentralized integral control is stable in all experiments, so that all four architectures can be compared on the same plant. The coupling sweep of Section 4.4 covers the range from zero to this nominal level and thus quantifies the dependence of the results on the coupling magnitude.

Slower and later cross-paths, Eqs. (13)–(14). Heat injected at zone i reaches zone j only after passing through the thermal capacitance of zone i and the lateral conduction path; the corresponding transfer function is therefore at least second order, with a dominant time constant not smaller than the slower of the two zones. When such a higher-order lag is approximated by an FOPDT model using the half rule, part of the neglected fast lag is added to the dominant time constant and part to the apparent delay [29]. The factor $\kappa_\tau = 1.2$ and the additional delay $\delta_d = 1$ sample represent this effect in its simplest form. The extra delay is referenced to the source heater because the heat originates there.

These rules reproduce the qualitative ordering expected from heat-transfer considerations, namely that cross-paths are weaker, slower, and later than direct paths, but the specific values are not calibrated to a particular chamber. The influence of the coupling magnitude is examined by the coupling sweep, and that of parameter errors by the mismatch scenarios of Section 4.4, in which the plant gains and time constants deviate by up to ±20% from the model; a separate variation of κ_τ, δ_d, and $\gamma_{\mathrm{D}}/\gamma_{\mathrm{N}}$ is left for future work (Section 5.5).

Remark 1 (Relation to Existing Multi-Channel TVAC Models). The use of interacting heater-to-zone transfer paths is motivated by earlier multi-channel TVAC modeling work. In particular, Li et al. [8] considered a six-channel heater and temperature-measurement configuration and characterized interactions among channels. The present study adopts the general control-oriented idea of a coupled MIMO thermal model, but does not reuse the experimentally identified parameters of that system. Instead, its deterministic direct-path parameters and topology-based cross-path coefficients define a synthetic test environment for controlled comparisons of decentralized and centralized architectures. Consequently, the results characterize the specified benchmark and should not be interpreted as experimental validation on a physical TVAC installation.

2.5 Unmodeled Physics and Validity Range

The linear time-invariant (LTI) MIMO-FOPDT plant neglects several effects that are important in thermal-vacuum testing. (i) *Radiation.* Radiative exchange with the shroud scales with $T^4 - T_{\mathrm{sh}}^4$; its linearized heat-transfer coefficient is $h_{\mathrm{r}} = \varepsilon\sigma(T + T_{\mathrm{sh}})(T^2 + T_{\mathrm{sh}}^2) \approx 4\varepsilon\sigma\bar{T}^3$ [30, 31]. Between -20 and $+80\,^{\circ}\mathrm{C}$, a range typical of TVAC cycling, $\bar{T}^3$ changes by a factor of $(353.15/253.15)^3 \approx 2.7$, so the effective loss conductance, and with it the local gains and time constants, varies by a comparable factor over a test cycle. (ii) *Temperature-dependent properties.* Thermal conductivity, specific heat, and contact conductance under vacuum vary with temperature and clamping conditions. (iii) *Boundary conditions.* The shroud temperature is itself a time-varying boundary rather than a constant ambient. (iv) *Actuators and sensors.* Heater power is proportional to the square of the applied voltage, and sensors add their own dynamics, offset, and noise. Consequently, the LTI benchmark is valid only in a neighborhood of one operating point; wide temperature cycles would require gain scheduling, linear-parameter-varying, or nonlinear prediction models [17, 18].

Because no experimental data are used, the results characterize the controllers within this local LTI description. The parametric mismatch scenarios of Section 4.4 expose the controllers to plant deviations of up to $\pm 20\%$ in gains and time constants, which is of the order of the operating-point dependence of the radiative loss over a moderate temperature range, but they do not reproduce the structural nonlinearity of radiation. Validation against a higher-fidelity nonlinear model or measured TVAC data is therefore identified as necessary future work (Section 5.5).

2.6 Discrete-Time Plant Representation

For the main nine-zone benchmark, the sampling interval is $T_s = 1\,\mathrm{s}$. Each active FOPDT path is simulated with an exact zero-order-hold discretization of its first-order lag and an integer input delay. Using the convention that $x_{ji}(k)$ is the state before the update at sample k, the path dynamics are

$$x_{ji}(k+1) = a_{ji}x_{ji}(k) + b_{ji}U_i(k - d_{ji}), \tag{15}$$

with

$$a_{ji} = \exp\left(-\frac{T_s}{\tau_{ji}}\right), \qquad b_{ji} = K_{ji}(1 - a_{ji}). \tag{16}$$

The aggregate temperature output is formed by superposing the active path contributions:

$$T_j(k) = T_{\mathrm{base},j} + \sum_{i:(j,i)\in\mathcal{E}} x_{ji}(k). \tag{17}$$

Here $T_{\mathrm{base},j}$ denotes the chosen temperature offset. The plant generator itself computes temperature increments relative to zero; an absolute-temperature simulation may add an ambient or nominal baseline separately. Equation (17) therefore distinguishes the incremental plant dynamics from the temperature offset used by the control experiments.

2.7 Control Objective and Operating Constraints

The objective is to track a prescribed temperature-reference vector $\mathbf{r}(k) \in \mathbb{R}^N$ while respecting actuator limits. For the main nine-zone experiments, the scaled heater commands satisfy

$$0 \le U_i(k) \le 5, \tag{18}$$

$$-0.25 \le \Delta U_i(k) \le 0.25, \tag{19}$$

where $\Delta U_i(k) = U_i(k) - U_i(k-1)$. These bounds are benchmark command limits, not experimentally established electrical power ratings. These limits are enforced directly in the constrained MPC optimization problems [13, 14]. For the PI-based controllers, the same actuator bounds are imposed through output saturation and slew-rate limiting, with anti-windup compensation.

TVAC test procedures also specify temperature limits for the test article and admissible temperature differences between regions, which protect hardware against overheating and thermally induced stress. These are expressed as

$$T_{\min} \le T_i(k) \le T_{\max}, \tag{20}$$

$$|T_i(k) - T_j(k)| \le \Delta T_{\max}, \qquad (i,j) \in \mathcal{E}_{\mathrm{N}}, \tag{21}$$

They are not imposed in the present controllers; their treatment is discussed in Sections 3.2.5 and 5.5. Controller-specific implementation details and the evaluation protocols are presented in the following sections.

3 Controller Design

Four control architectures and one ablation variant are considered to assess the effects of cross-channel model information, predictive coordination, and constraint handling: decentralized PI, decentralized MPC, static-decoupled MIMO-PI, and centralized constrained MIMO-MPC. The first three serve as reference controllers, whereas coupling-aware MIMO-MPC is the principal architecture investigated in this study. All controllers act on the same coupled thermal plant and use the sampling period and actuator limits specified in Section 2. For the primary benchmark, $N = 9$; the predictive formulation below is stated for a general N-zone plant. The controllers differ in the model information available to their predictions or compensators, not in the physical plant on which their commands are evaluated.

3.1 Reference Control Architectures

3.1.1 Decentralized PI

Each zone is controlled by an independent PI loop tuned from its direct FOPDT path $G_{ii}(s)$. With $e_i(k) = r_i(k) - T_i(k)$, the nominal unconstrained PI law is

$$U_{i,\mathrm{raw}}^{\mathrm{PI}}(k) = K_{p,i}\left[e_i(k) + \frac{T_s}{T_{I,i}}\sum_{\ell=0}^{k} e_i(\ell)\right], \tag{22}$$

where the summation denotes the ideal integral action before anti-windup correction. The controller gains follow the SIMC-style rule [29]:

$$K_{p,i} = \frac{\tau_{ii}}{K_{ii}(\theta_{ii} + \lambda_i)}, \qquad T_{I,i} = \min\{\tau_{ii}, 4(\theta_{ii} + \lambda_i)\}, \tag{23}$$

where $\lambda_i = \max(\lambda, T_s)$ and $\lambda = 20$ s in the primary benchmark. The integral state is updated in discrete time and the resulting unconstrained command is subject to amplitude saturation and sample-to-sample slew-rate limiting. Back-calculation anti-windup then corrects the integral state using the difference between the unconstrained command and the actually applied command; the back-calculation coefficient is $K_{\mathrm{aw}} = 0.25$ in the primary benchmark. Thus, the applied heater command satisfies the same actuator limits as the MPC controllers. The nine PI loops do not explicitly compensate for off-diagonal thermal interactions.

3.1.2 Decentralized MPC

Each heater is controlled by a local, constrained step-response MPC using only its direct dynamic path $G_{ii}(s)$. Let $\boldsymbol{\phi}_i \in \mathbb{R}^{N_p \times N_c}$ be the lower-triangular prediction matrix constructed from direct-path step-response coefficients, and let $\mathbf{y}_{0,i}$ be the corresponding free response. The local controller solves

$$\min_{\Delta \mathbf{u}_i} \; \|\mathbf{y}_{0,i} + \boldsymbol{\phi}_i \Delta \mathbf{u}_i - \mathbf{r}_i\|_2^2 + \lambda_u \|\Delta \mathbf{u}_i\|_2^2, \tag{24}$$

subject to the same input-amplitude and input-rate limits as centralized MPC. Here $\mathbf{r}_i \in \mathbb{R}^{N_p}$ contains the prescribed future reference samples for zone i.

The local free response is propagated from the direct-path model and the previously applied heater-input history. A constant output correction equal to the difference between the measured temperature deviation and the corresponding local-model output is then added across the prediction horizon. This accounts for the current effect of unmodeled contributions, including cross-zone heating, without incorporating off-diagonal dynamic paths into $\boldsymbol{\phi}_i$. The local QPs are solved independently, although their commands act simultaneously on the coupled plant. Both MPC architectures use the prescribed future reference trajectory over the prediction horizon. The step-response construction and receding-horizon procedure are detailed for the centralized case in Section 3.2.

3.1.3 Static-Decoupled MIMO-PI

The static-decoupled MIMO-PI baseline uses the complete steady-state gain matrix $\mathbf{K} = \mathbf{G}_N(0)$, but not the full cross-channel dynamics. A damped inverse defines the input transformation [28]:

$$\mathbf{D}_{\text{dec}} = \mathbf{K}^\top \left(\mathbf{K}\mathbf{K}^\top + \mu^2 \mathbf{I}_N\right)^{-1}, \qquad \mu = 0.05. \tag{25}$$

The virtual PI commands $\mathbf{v}(k)$ are mapped to unconstrained physical heater commands by $\mathbf{U}_{\text{raw}}(k) = \mathbf{D}_{\text{dec}}\mathbf{v}(k)$. The virtual loops are tuned using the diagonal entries of the effective gain matrix $\mathbf{K}_{\text{eff}} = \mathbf{K}\mathbf{D}_{\text{dec}}$ and the original direct-channel time constants and delays, following Eq. (23). The same $\lambda = 20$ s is used as for decentralized PI in the primary benchmark. Physical heater commands are subject to amplitude saturation and slew-rate limiting. Back-calculation anti-windup uses the discrepancy between the unconstrained and applied physical commands, mapped through $\mathbf{K}$ to correct the virtual-loop integral states; its coefficient is $K_{\text{aw}} = 0.25$. This baseline compensates for steady-state interaction but does not predict heterogeneous cross-channel time constants and delays.

3.2 Coupling-Aware Centralized Constrained MIMO-MPC

3.2.1 Operating Principle and Closed-Loop Architecture

The centralized MPC coordinates all N heater inputs using the complete MIMO-FOPDT model. At each sampling instant, the controller constructs a free thermal response from its current prediction state and the previously applied heater-input history, predicts temperature deviations over N_p steps, optimizes a sequence of N_c future heater increments, and applies only the first optimized increment. The prediction and optimization are repeated at the next sampling instant. The state information used to construct the free response is specified in Section 3.2.3. The receding-horizon formulation follows standard MPC principles [13, 14].

Fig. 3 summarizes the conceptual closed-loop architecture. The coupled model predicts how each heater influences all dynamically connected zones, and the constrained optimizer computes coordinated heater increments. The diagram depicts the nine-zone benchmark and the input constraints included in the baseline QP; no temperature-output constraints are imposed in the optimization.

3.2.2 Cross-Coupled Step-Response Prediction

For each active heater-to-zone path (j, i), define the discrete-time FOPDT coefficients

$$a_{ji} = \exp\left(-\frac{T_s}{\tau_{ji}}\right), \qquad b_{ji} = K_{ji}(1 - a_{ji}), \qquad d_{ji} = \frac{\theta_{ji}}{T_s}, \tag{26}$$

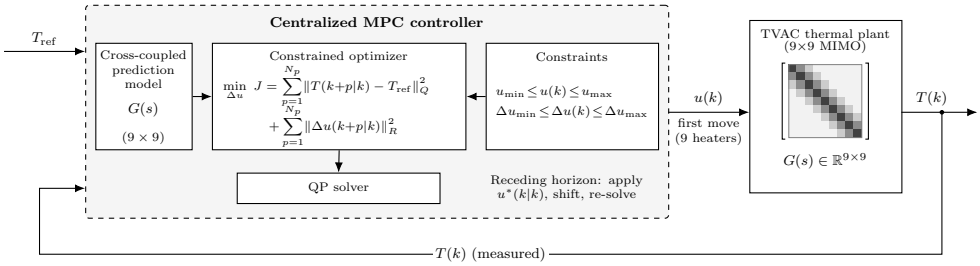

Fig. 3 Conceptual closed-loop architecture of coupling-aware centralized MIMO-MPC for the nine-zone synthetic benchmark. The full cross-coupled model is used to optimize heater increments subject to input-amplitude and input-rate constraints. Only the first optimized control move, comprising nine heater increments, is applied before the optimization is repeated.

where d_{ji} is a nonnegative integer for the main benchmark. The step-response coefficient at prediction step p is

$$S_{ji}(p) = \begin{cases} K_{ji}\left(1 - a_{ji}^{p-d_{ji}}\right), & p > d_{ji}, \\ 0, & p \le d_{ji}. \end{cases} \tag{27}$$

Inactive paths have $S_{ji}(p) = 0$. For each p, collect the channel responses into $\mathbf{S}_p = [S_{ji}(p)]_{j,i=1}^{N}$ and set $\mathbf{S}_p = \mathbf{0}$ for $p \le 0$. The dynamic prediction matrix is

$$\mathbf{\Phi} = \begin{bmatrix} \mathbf{S}_1 & \mathbf{0} & \cdots & \mathbf{0} \\ \mathbf{S}_2 & \mathbf{S}_1 & \cdots & \mathbf{0} \\ \vdots & \vdots & \ddots & \vdots \\ \mathbf{S}_{N_p} & \mathbf{S}_{N_p-1} & \cdots & \mathbf{S}_{N_p-N_c+1} \end{bmatrix} \in \mathbb{R}^{NN_p \times NN_c}. \tag{28}$$

The off-diagonal entries of $\mathbf{S}_p$ encode the magnitudes, response speeds, and delays of spatial interactions. Retaining these entries distinguishes the centralized coupled prediction from the local direct-path predictions used by decentralized MPC. The step-response (dynamic-matrix) construction follows standard DMC formulations [13, 16]; the study evaluates the consequences of retaining full dynamic coupling information rather than proposing a new prediction algorithm.

Define the stacked input-increment vector

$$\Delta\mathcal{U}_k = \left[\Delta\mathbf{U}(k|k)^\top \ \cdots \ \Delta\mathbf{U}(k+N_c-1|k)^\top\right]^\top. \tag{29}$$

The predicted temperature deviations satisfy

$$\mathcal{Y}_k = \mathcal{Y}_{0,k} + \mathbf{\Phi}\Delta\mathcal{U}_k, \tag{30}$$

where $\mathcal{Y}_{0,k} \in \mathbb{R}^{NN_p}$ is the free response associated with the current predictor states, previously applied heater inputs, and their delayed effects.

3.2.3 State Information and Output-Bias Correction

Because the information used to construct $\mathcal{Y}_{0,k}$ can strongly influence the comparison of architectures, two information settings are distinguished and are stated for every experiment (Table 3).

Simulator-state prediction (setting F). The free response is obtained by propagating the simulated path states $x_{ji}(k)$ and input-delay buffers of the plant under held heater commands. This setting is an idealized upper bound on prediction quality; it is not a state estimator realizable on a physical TVAC installation. It is used for MIMO-MPC in the nominal, noise-free implementations, including the four-controller comparison of Study A.

Measurement-based prediction (setting M). The MPC propagates its internal model from the applied heater-input history and adds a constant output-bias correction over the prediction horizon,

$$\hat{\mathbf{d}}(k) = \mathbf{y}_{\mathrm{m}}(k) - \hat{\mathbf{y}}(k), \qquad \mathcal{Y}_{0,k} \leftarrow \mathcal{Y}_{0,k} + \mathbf{1}_{N_p} \otimes \hat{\mathbf{d}}(k), \tag{31}$$

where $\mathbf{y}_{\mathrm{m}}(k)$ is the measured temperature deviation and $\hat{\mathbf{y}}(k)$ the internal-model output. This is the step-disturbance model of offset-free MPC [32–34]; it does not modify $\mathbf{\Phi}$ and is an algebraic output correction rather than a separately validated disturbance-state observer. Decentralized MPC always uses this setting with its local model, and MIMO-MPC uses it with the full model in the designated uncertainty and robustness implementations, which include the perturbed protocol used for the matched-setting coupling sweep and the scalability tests (Section 4). When both MPCs operate in setting M with identical settings, they receive the same measurements and correction law and differ only in their prediction matrices.

Noise transmission. Because the correction uses the current measurement directly, each noise sample $\mathbf{v}(k)$ shifts the entire predicted trajectory. Ignoring active constraints and the small regularization ε introduced below, the first optimal move then contains the term $-\mathbf{G}_v\mathbf{v}(k)$, with

$$\mathbf{G}_v = \begin{bmatrix}\mathbf{I}_N & \mathbf{0}\end{bmatrix}\left(\mathbf{\Phi}^\top\mathbf{\Phi} + \lambda_u\mathbf{I}\right)^{-1}\mathbf{\Phi}^\top(\mathbf{1}_{N_p} \otimes \mathbf{I}_N). \tag{32}$$

The noise-to-input gain $\mathbf{G}_v$ increases as the move penalty λ_u decreases; it is diagonal for decentralized MPC but a full matrix for MIMO-MPC, so noise measured in one zone also moves the heaters of neighboring zones. This relation is used in Section 5.3 to interpret the noise sensitivity of the tuned controllers.

3.2.4 Tracking Objective and Quadratic Program

Let $\mathcal{R}_k \in \mathbb{R}^{NN_p}$ denote the stacked reference deviations over the prediction horizon. The controller minimizes

$$J_k = \|\mathcal{Y}_k - \mathcal{R}_k\|^2_{\mathbf{Q}} + \|\Delta\mathcal{U}_k\|^2_{\mathbf{R}}\,, \tag{33}$$

where $\|\mathbf{z}\|_{\mathbf{W}}^2 = \mathbf{z}^\top \mathbf{W}\mathbf{z}$, $\mathbf{Q} = \mathbf{I}_{NN_p}$, and $\mathbf{R} = \lambda_u \mathbf{I}_{NN_c}$ with $\lambda_u > 0$. Substituting Eq. (30) yields the quadratic program

$$\min_{\Delta\mathcal{U}_k} \frac{1}{2}\Delta\mathcal{U}_k^\top \mathbf{H}\Delta\mathcal{U}_k + \mathbf{f}_k^\top \Delta\mathcal{U}_k, \tag{34}$$

with

$$\mathbf{H} = 2\left(\mathbf{\Phi}^\top \mathbf{Q}\mathbf{\Phi} + \mathbf{R}\right) + \varepsilon \mathbf{I}_{NN_c}, \tag{35}$$

$$\mathbf{f}_k = 2\mathbf{\Phi}^\top \mathbf{Q}\left(\mathcal{Y}_{0,k} - \mathcal{R}_k\right), \tag{36}$$

where $\varepsilon = 10^{-8}$ is a numerical regularization term added to the Hessian after multiplication by two, as in the implementation. For $\lambda_u > 0$, the Hessian is positive definite, and the QP has a unique optimizer whenever its linear constraint set is nonempty.

3.2.5 Actuator Constraints and Receding-Horizon Update

For the primary nine-zone benchmark, the predicted input increments satisfy

$$-0.25\,\mathbf{1}_{NN_c} \le \Delta\mathcal{U}_k \le 0.25\,\mathbf{1}_{NN_c}. \tag{37}$$

Define the lower-triangular accumulation matrix

$$\mathbf{L} = \mathrm{tril}(\mathbf{1}_{N_c \times N_c}) \otimes \mathbf{I}_N. \tag{38}$$

The stacked future absolute heater commands are

$$\mathcal{U}_k = (\mathbf{1}_{N_c} \otimes \mathbf{I}_N)\mathbf{U}(k-1) + \mathbf{L}\Delta\mathcal{U}_k, \tag{39}$$

and must satisfy

$$\mathbf{0} \le \mathcal{U}_k \le 5\,\mathbf{1}_{NN_c}. \tag{40}$$

After solving the constrained problem, only the first optimal increment is implemented:

$$\mathbf{U}(k) = \mathbf{U}(k-1) + \Delta\mathbf{U}^\star(k|k). \tag{41}$$

The remaining planned moves are discarded and the optimization is repeated at $k+1$.

Feasibility. Because the applied input always satisfies $\mathbf{0} \le \mathbf{U}(k-1) \le 5\,\mathbf{1}_N$, the zero increment $\Delta\mathcal{U}_k = \mathbf{0}$ satisfies Eqs. (37) and (40) at every sample, provided that the initial input is admissible. The input-constrained QP is therefore always feasible and, since $\mathbf{H} \succ 0$, has a unique solution. A solver status other than success can thus arise only from numerical causes, such as an iteration limit or ill-conditioning, and not from infeasibility. The implementation checks the solver exit flag and applies the zero increment, which is always admissible, if a valid solution is not returned. For the primary nine-zone comparison, the nominal predictive settings are $N_p = 80$, $N_c = 20$, and $\lambda_u = 0.05$. Experiments involving controller tuning, tightened actuator limits, or larger arrays use the settings reported in their respective evaluation protocols.

Scope of the constraints. The QP constrains actuator commands only. The temperature limits of Eqs. (20)–(21) are linear in the predicted outputs and could be added to Eq. (34); to avoid infeasibility under disturbances and mismatch [35], they would be imposed as soft constraints with slack variables and an exact penalty [36], which preserves convexity and guarantees feasibility. Such constraints, which are coupled across zones for Eq. (21) and can therefore be anticipated only by a joint prediction, are not included in the present evaluation (Section 5.5).

3.2.6 Coordination Without Coupling Information

If the prediction matrix is block diagonal after permutation, that is, if only direct paths are modeled, then with $\mathbf{Q}$ and $\mathbf{R}$ diagonal and the channel-wise constraints of Eqs. (37)–(40), the Hessian, the gradient, and the feasible set of the centralized QP are all separable by heater. The joint problem then decomposes exactly into the N local problems of Eq. (24), and with the measurement-based correction of Eq. (31) its solution coincides with that of decentralized MPC. Joint optimization alone therefore provides no benefit: under matched information and settings, any difference between the two architectures arises from the coupling terms of the prediction model. This property is used as a consistency check in the coupling sweep of Section 5.3.

4 Simulation Design

The simulations test whether modeling dynamic cross-zone interaction improves constrained temperature tracking relative to local prediction and static decoupling. Four experiments (Table 3) are run: (A) a common-condition comparison of four controllers, (B) independently tuned MPC evaluation with actuator-constraint tests, (C) coupling sensitivity and robustness, and (D) validation across array sizes. All plants are synthetic; the experiments are not a validation on physical TVAC hardware.

4.1 Common Setup

The nominal nine-zone benchmark is defined in Section 2. Within each study, controllers act on the same coupled plant (ambient 25 °C, $0 \le U_i \le 5$, $|\Delta U_i| \le 0.25$ per sample) and, where stochastic, share noise and perturbation realizations. The two MPCs share reference preview and solver (`quadprog`), and share horizons and move penalty in the matched studies. The PI controllers use the tuning and anti-windup rules of Section 3 and receive no future reference; this information difference is retained as a property of the architectures. Because durations, sampling periods, and reference amplitudes differ between studies, errors are compared within each protocol only.

4.2 Study A: Nominal Four-Controller Comparison

The reference follows from a feasible steady-state heater vector,

$$\mathbf{u}^\star = \mathrm{linspace}(1.5, 2.8, 9)^\top, \qquad \Delta\mathbf{T}^\star = \mathbf{K}\mathbf{u}^\star, \tag{42}$$

Table 2 Model information, state information, and coordination in the evaluated controllers.

Controller	Model information	State information	Prediction	Input handling
Decentralized PI	Direct-path tuning	Measured output	None	Limited, anti-windup
Decentralized MPC	Direct dynamics	Setting M	Local	Constrained QP
Static-decoupled MIMO-PI	Full static gain	Measured output	None	Limited, anti-windup
Centralized MIMO-MPC	Full dynamic MIMO	F (nominal), M (perturbed)	Joint	Constrained QP

Table 3 Simulation protocols. Settings within a study are shared by the controllers being compared unless stated otherwise. "Matched" denotes identical (N_p, N_c, λ_u) for both MPCs; "tuned" denotes separately selected settings. F and M denote the state-information settings of Section 3.2.3.

Study	Main question	T_s/T_{sim} (s)	MPC settings (N_p, N_c, λ_u)	Evaluation conditions
A	Does full dynamic coupling improve tracking against three baselines?	1/4000	Matched $(80, 20, 0.05)$	Four controllers; nine zones; four-stage reference; nominal conditions; MIMO-MPC in setting F.
B	Does the advantage persist after independent MPC tuning?	1/1000 (tuning), 1/1600 (test)	Tuned on a common 36-point grid	Decentralized vs. centralized MPC; separate tuning and held-out trajectories; constrained QP vs. clipping.
C	How do coupling and disturbances affect tracking?	1/1600	Tuned (Study B) ; matched sweep in D protocol	Five coupling factors; five robustness scenarios; 30 paired trials for stochastic scenarios.
D	How does performance change with plant dimension?	2/400	Matched $(30, 8, 0.05)$	Four controllers; 9, 16, 25, 36 zones; perturbed, setting M; coupling and robustness checks.

so every zone has its own target; the zone targets therefore differ, although all zones move in the same direction. The 4000-s profile has four conditions: a ramp to 75% of the target deviation (0–800 s), a hold (800–1600 s), a ramp to the full target (1600–2400 s), and a sinusoid about the target (2400–4000 s; amplitude 0.6 °C, period 1600 s). Decentralized PI, decentralized MPC, static-decoupled MIMO-PI, and coupling-aware centralized MIMO-MPC are compared under identical nominal conditions.

4.3 Study B: MPC Tuning and Constraints

Both MPCs are searched over the same grid,

$$\begin{aligned} N_p &\in \{40, 80, 120\}, \quad N_c \in \{10, 20, 30\}, \\ \lambda_u &\in \{0.005, 0.05, 0.5, 2\}, \end{aligned} \tag{43}$$

and scored by RMSE on a 1000-s tuning trajectory. The grid is chosen as follows. With dominant time constants of 68–116 s and delays of 2–6 s, $N_p = 40$, 80, and 120 samples cover approximately 0.35–0.6, 0.7–1.2, and 1.0–1.8 dominant time constants, that is, horizons from shorter than to longer than the main transient. $N_c \in \{10, 20, 30\}$ spans a quarter to three quarters of the shortest N_p, and the move penalty spans almost three decades, from aggressive (0.005) to strongly damped (2) action. The selection criterion is the minimum RMSE on the noise-free tuning trajectory, without further refinement. Each architecture's best setting is then tested on a distinct 1600-s trajectory with ramp, hold, ramp, and sinusoid phases of 320, 320, 320, and 640 s (sinusoid amplitude 0.35 °C), which is not used for tuning and therefore measures generalization to a different noise-free trajectory. A separate test compares the centralized MPC's constrained QP with an unconstrained solution followed by actuator clipping, using the same model and tuning, to isolate the value of enforcing limits inside the optimization. Because the two architectures select different settings, the results of Studies B and C are labeled "tuned" and are not used to infer the effect of coupling information; that effect is assessed with matched settings in Studies A, C (matched sweep), and D.

4.4 Study C: Coupling Sensitivity and Robustness

Coupling is varied by scaling the off-diagonal gains,

$$\begin{aligned} \mathbf{K}(\gamma) &= \operatorname{diag}(\mathbf{K}) + \gamma\big[\mathbf{K} - \operatorname{diag}(\mathbf{K})\big], \\ \gamma &\in \{0, 0.25, 0.5, 0.75, 1\}, \end{aligned} \tag{44}$$

with a fixed target $\Delta\mathbf{T}^\star = 0.5\,\operatorname{diag}(\mathbf{K})$ (feasibility checked at each γ) and both MPC models rebuilt for each plant. This tests the effect of *modeled* coupling strength, not robustness to an unannounced coupling change. Two sweeps are reported: one with the tuned settings of Study B, and one with matched settings in the perturbed Study-D protocol, in which both MPCs use the measurement-based correction of Eq. (31).

Robustness uses sensor noise, an external disturbance, parametric mismatch, and their combination (Table 4). The disturbance is added to zones 2, 5, and 8 for $0.45T_{\text{sim}} \le t < 0.65T_{\text{sim}}$. Mismatch multiplies the gain and time constant of each

Table 4 Robustness scenarios for Study C.

Scenario	Noise SD [°C]	Load [°C]	Mismatch ρ
S1 Nominal	0	0	0
S2 Noise	0.05	0	0
S3 Load	0	0.15	0
S4 Mismatch	0	0	0.10
S5 Combined	0.05	0.15	0.20

active plant channel by $1 + \rho\xi$, $\xi \sim \mathcal{U}(-1, 1)$, with delays unchanged; controllers keep their nominal models. Stochastic scenarios use 30 paired realizations (common seeds for both MPCs); the nominal and deterministic-load cases need no random draws.

The scenario levels are chosen as follows. A noise standard deviation of 0.05 °C is used as a moderate random-noise level; it is smaller than the tolerance of a Class A platinum resistance thermometer, $\pm(0.15 + 0.002|t|)$ °C [37], and therefore represents random fluctuation rather than calibration error. The load of 0.15 °C equals approximately 24–34% of the per-zone target excursion $0.5K_{ii}$ used in Study C and represents an unmodeled heat input, such as dissipation of the test article, of the same order as the commanded change. Parametric mismatch of ±10% in gains and time constants represents a moderate identification error, and ±20% in the combined scenario is a stress case. The mismatch scenarios also serve as a partial sensitivity analysis of the plant parameters, because the plant deviates from the model while the controllers retain the nominal parameter rules.

4.5 Study D: Plant Dimension

Square grids of 9, 16, 25, and 36 heater–zone pairs are simulated with $T_s = 2$ s, $T_{\text{sim}} = 400$ s, and $(N_p, N_c, \lambda_u) = (30, 8, 0.05)$, with sensor noise 0.05 °C, load 0.15 °C, and mismatch bound 0.20 applied together. The nine-zone case keeps the original gains and time constants (delays converted to the new sampling interval; off-diagonal Frobenius ratio 0.406). Larger plants are synthetic: direct-channel parameters are interpolated over the grid, edge-adjacent zones are coupled, and the off-diagonal Frobenius ratio is normalized to 0.30. Consequently, the nine-zone plant and the larger plants differ in their coupling rule (with and without diagonal neighbors) and coupling level, so the study compares controllers within each size and does not isolate the effect of dimension alone (Section 5.5). The study examines how accuracy, optimization size, and solution time vary with the number of zones; no experimental scalability is claimed.

The centralized QP has NN_c decision variables with a dense Hessian, $2NN_c$ simple bounds on the increments, and $2NN_c$ linear inequalities for the amplitude limits; each local QP has N_c variables, $2N_c$ bounds, and $2N_c$ inequalities. With $N_c = 8$, the centralized problem has 72, 128, 200, and 288 decision variables for $N = 9$, 16, 25, and 36, whereas each of the N local problems has 8. For decentralized MPC, the reported time is the sum of the N local solution times per step, which corresponds to sequential execution on one processor.

4.6 Performance Measures and Reproducibility

For N_t samples and n zones, with $e_i(k) = r_i(k) - T_i(k)$, the global RMSE is

$$\mathrm{RMSE} = \sqrt{\frac{1}{nN_t} \sum_{k=1}^{N_t} \sum_{i=1}^{n} e_i^2(k)}, \tag{45}$$

and the mean IAE, mean total input variation (TV), and maximum error are

$$\mathrm{MeanIAE} = \frac{1}{n} \sum_{i=1}^{n} \int_0^{T_{\mathrm{sim}}} |e_i(t)|\, dt, \tag{46}$$

$$\mathrm{MeanTV} = \frac{1}{n} \sum_{i=1}^{n} \sum_{k=2}^{N_t} |U_i(k) - U_i(k-1)|, \tag{47}$$

$$e_{\max} = \max_{i,k} |e_i(k)|. \tag{48}$$

Errors are computed from the true plant temperatures, not from the noisy measurements. Zone-wise and condition-wise RMSE, compliance with magnitude and rate limits, and, for MPC, QP success rate and solution time are also reported. Stochastic trials use fixed seeds and paired realizations. Uncertain cases perturb the physical plant or measurements without informing the controllers.

Implementation details. The delay buffers and the internal models of the MPCs are initialized consistently with this steady state. Measurement noise is zero-mean, independent, identically distributed Gaussian noise added to the measured outputs only. The prediction matrices are constructed as in Eqs. (27)–(28). The QPs are solved with in MATLAB 2025b on Intel Core Ultra 7 RAM 32 GB. Computation times characterize this implementation, not real-time guarantees on embedded hardware.

5 Results and Discussion

Results follow the four studies of Section 4; comparisons are made within a protocol, and all plants are synthetic.

5.1 Nominal Comparison of Control Architectures

Table 5 and Fig. 4 summarize Study A. The coupling-aware centralized MIMO-MPC gives the lowest global RMSE (0.0008 °C) against 0.0097 (decentralized MPC), 0.0265 (decentralized PI), and 0.0489 °C (static-decoupled MIMO-PI), together with the lowest IAE and maximum error, and the same ordering holds in all four conditions. Its input variation (TV 4.28) is close to decentralized PI (4.11) and below decentralized MPC (5.01), and no amplitude or rate limit is reached.

The roughly 12-fold RMSE reduction relative to decentralized MPC is not an isolated coupling effect. In this nominal run MIMO-MPC builds its free response from the simulated full-model states and delayed-input history (setting F), whereas decentralized MPC uses direct-path prediction with a measured-output offset correction; the

gap therefore reflects both cross-zone information and prediction-state initialization. The matched comparisons of Sections 5.3 and 5.4, in which both MPCs use setting M, give a much smaller advantage of MIMO-MPC: 12.8% at nominal coupling in the matched sweep and a factor of 1.9 in the perturbed nine-zone case of Study D. The large ratio of Study A should therefore be attributed predominantly to the information setting and to the large, unperturbed ramps, and only partly to coupling information.

Lag analysis of the tracking error. Because the temperature traces of the MPCs in Fig. 4 are nearly indistinguishable from the reference, the tracking error is characterized by its ramp-phase lag. For decentralized PI, MIMO-PI, and decentralized MPC, the ratio of the C1 and C3 RMSE equals 3.0, the ratio of the two ramp rates (0.75/0.25), so their ramp error is essentially a constant time lag $\tau_{\mathrm{lag}} = \bar{e}/\dot{r}$. Averaged over zones, τ_{lag} is 11.8 s for decentralized PI, 22.1 s for MIMO-PI, -4.2 s for decentralized MPC, and -0.05 s for MIMO-MPC, where a negative lag means that the output leads the reference. Both MPCs begin to move their inputs about $N_p = 80$ samples before the C3 ramp starts at $t = 1600$ s because of reference preview, whereas the PI controllers react only after the reference changes. The lead of decentralized MPC is consistent with its local model assuming the direct gain K_{ii}, whereas the effective common-mode gain is about 1.9 times larger (see below); only MIMO-MPC combines preview with the correct gain, which brings its lag close to zero. The inputs remain inside the limits throughout (peak input 3.31; peak $|\Delta u_i|$ 0.129 and 0.162 for decentralized MPC and MIMO-MPC, and 0.022 and 0.015 for the two PI controllers, against a limit of 0.25), so no constraint is active in Study A.

The advantage over PI likewise reflects reference preview and tuning. Static decoupling does not improve PI tracking. Because a reference that raises all zones together also drives neighboring heaters, each decentralized PI loop experiences an effective gain $\sum_j K_{ij}/K_{ii}$ of 1.88 on average (range 1.62–2.34) and therefore responds about 1.9 times faster than designed, consistent with the measured lag ratio $22.1/11.8 = 1.87$. Static decoupling removes this beneficial common-mode interaction and restores the designed closed-loop time constant ($\lambda = 20$ s plus delay), which explains its slower ramp tracking. This interpretation is specific to same-direction references; for references in which neighboring zones must move in opposite directions, interaction would oppose rather than assist the decentralized loops (Section 5.5).

Interpretation of the error magnitude. An RMSE of 0.0008 °C (0.8 mK) is about 250 times smaller than the tolerance of a Class A platinum resistance thermometer at 25 °C (±0.20 °C [37]) and below the noise and resolution of typical temperature acquisition chains. It is achievable only because Study A uses an exact, noise-free LTI plant and simulator-state prediction. These values should therefore be read as the error floor of an idealized benchmark for comparing architectures, not as achievable accuracy on a physical TVAC; on hardware, sensor accuracy, noise, and the unmodeled dynamics of Section 2.5 would determine the attainable error, as the noise scenarios of Section 5.3 illustrate.

5.2 MPC Tuning and Actuator Constraints

The selected settings (N_p, N_c, λ_u) are $(80, 10, 0.005)$ for decentralized MPC and $(40, 30, 0.005)$ for MIMO-MPC, with held-out RMSE of 0.0154 and 0.0024 °C. Because

Table 5 Study A: nominal tracking performance ($T_{\text{sim}} = 4000$ s, $N_p = 80$, $N_c = 20$, $\lambda_u = 0.05$; MIMO-MPC in setting F). IAE and TV follow Section 4.6.

Controller	RMSE [°C]	IAE	Max $\lvert e \rvert$ [°C]	TV	Condition RMSE [°C]			
					C1 Ramp	C2 Hold	C3 Ramp	C4 Sine
Dec. PI	0.0265	81.0	0.0742	4.11	0.0476	0.0047	0.0159	0.0220
Dec. MPC	0.0097	27.7	0.0416	5.01	0.0181	0.0012	0.0060	0.0074
MIMO-PI	0.0489	149.4	0.1353	4.05	0.0903	0.0102	0.0301	0.0375
MIMO-MPC	**0.0008**	**1.43**	**0.0222**	4.28	**0.0016**	**0.0003**	**0.0004**	**0.0005**

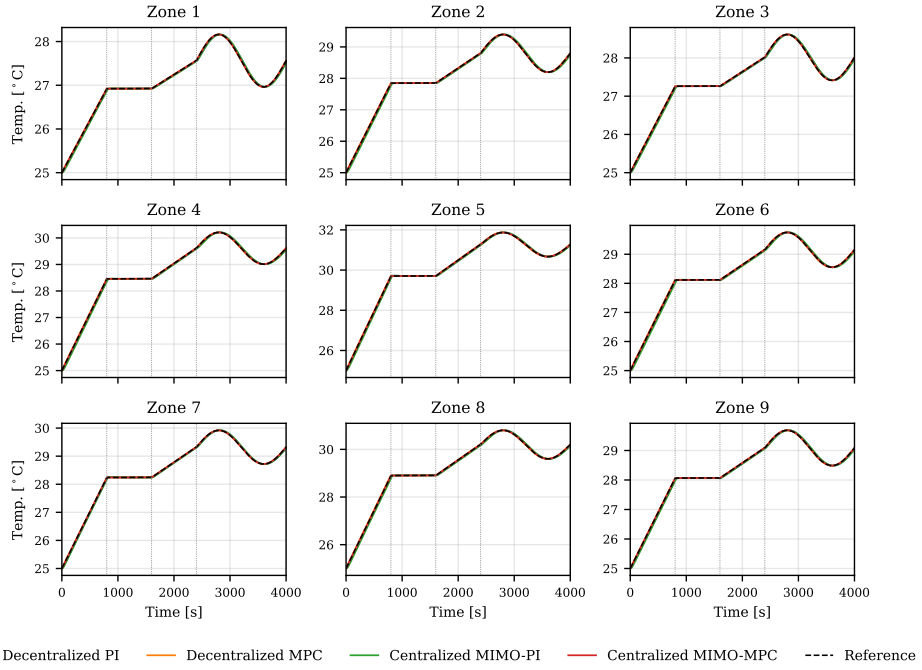

Fig. 4 Study A: temperature tracking in all nine zones under the four-stage reference; vertical lines mark condition boundaries. The MPC traces are nearly indistinguishable from the reference at this scale; the ramp-phase lags and input levels are quantified in Section 5.1.

the settings differ substantially, this result reflects the selection procedure, not the isolated effect of cross-zone modeling; the coupling analysis below addresses that question.

Sensitivity to the settings. Over the 36-point grid, decentralized MPC is insensitive to its settings (worst-to-best RMSE ratio 3.6), whereas MIMO-MPC is highly sensitive (ratio 21.8). MIMO-MPC prefers a short prediction horizon with many free moves: its grid-mean RMSE decreases from 0.0139 to 0.0041 °C as N_c increases from 10 to 30 and increases from 0.0037 to 0.0146 °C as N_p increases from 40 to 120. The tuned gap of about six between the two MPCs is therefore attributable mainly to the selected horizons. Both selected settings use the smallest move penalty, $\lambda_u = 0.005$, because the

tuning trajectory is noise free; the selected settings generalize to the held-out noise-free trajectory, but Section 5.3 shows that they do not generalize well to noisy operation. Results with separately selected settings (Studies B and C) and with matched settings (Studies A, C matched sweep, and D) are accordingly interpreted separately.

Tightening the amplitude bound to $u_{\max} = 3$ raises the tuned MIMO-MPC RMSE from 0.0024 to 0.0100 °C and activates the bound in 2.0% of the samples, and tightening the rate bound to $|\Delta u_i| \leq 0.05$ gives 0.0093 °C. Including the limits in the QP gives lower RMSE than solving without them and clipping, but only modestly (about 5% and 1.5%). In the tight-rate case, the solver did not return a successful status in two of 1600 steps, at which the zero increment was applied. Because the input-constrained QP is always feasible (Section 3.2.5), these events are not caused by infeasibility of the optimization problem but by numerical termination of the solver, for example at its iteration limit, in a problem with many simultaneously active rate constraints and a small move penalty ($N_c = 30$, $\lambda_u = 0.005$).The results thus support constraint feasibility and predictable fallback behavior rather than a large tracking benefit from QP constraint handling.

5.3 Coupling Sensitivity and Robustness

The nominal off-diagonal Frobenius ratio is 0.406, indicating appreciable interaction. Two sweeps are reported (Fig. 5): independently tuned MPCs, and matched settings in the perturbed Study D protocol. They differ in settings and references and are interpreted separately.

In the matched sweep the two MPCs coincide at $\gamma = 0$ (RMSE 0.038884 °C), as predicted by the decomposition property of Section 3.2.6; the dense centralized QP agrees with the decomposed solution to 3.3×10^{-8} °C. This is a consistency check, not evidence of better control. As coupling increases, the centralized controller's advantage grows progressively, to 2.2, 5.7, 9.1, and 12.8% for $\gamma = 0.25$, 0.5, 0.75, and 1, because the MIMO-MPC RMSE decreases by 11.4% (neighboring heaters provide additional authority) while that of decentralized MPC increases by 1.7%. Because both MPCs use the same measurements, correction law, horizons, and weights in this sweep, and joint optimization without coupling terms yields no benefit, this growth isolates the contribution of coupling information. This shows directly that the benefit of modeling interactions depends on their strength. The static-decoupled MIMO-PI, which uses only steady-state coupling information, equals decentralized PI at $\gamma = 0$ but is 1.66 times worse at $\gamma = 1$. Static interaction information applied through a decoupler thus does not reproduce the benefit obtained with dynamic coupling information in the predictor, although this comparison also involves the difference between PI and MPC. By contrast, the independently tuned MPCs keep a large gap already at $\gamma = 0$, where coupling is absent; that gap reflects the selected horizons and move penalties and is not an architectural benefit.

Table 6 and Fig. 6 give the robustness results. MIMO-MPC has the lower mean RMSE in all scenarios, but its nominal advantage shrinks with noise: the decentralized-to-centralized ratio falls from 6.32 (nominal) to 1.41 (noise) and 1.36 (combined), whereas mismatch alone has a smaller effect (5.62). Relative to the nominal case, the MIMO-MPC RMSE grows by factors of 8.8 (noise), 3.7 (load), 1.13 (mismatch), and

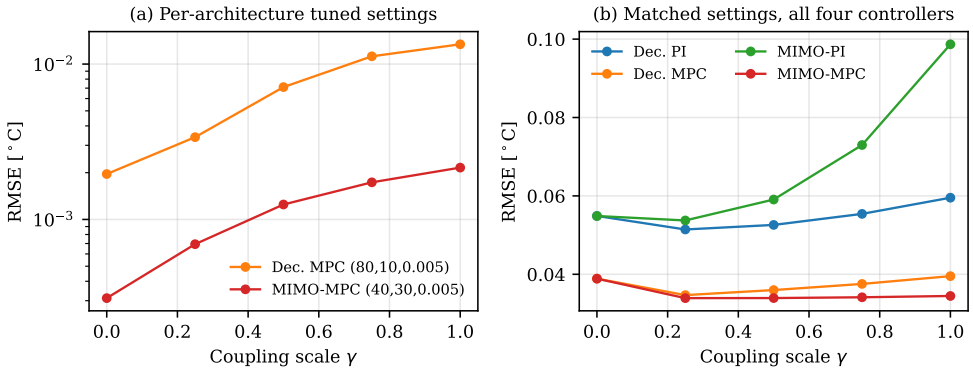

Fig. 5 Tracking RMSE versus coupling scale: (a) independently tuned MPCs (Study C); (b) four controllers with matched settings in the perturbed Study D protocol. The panels use different settings and conditions and are not a single controlled experiment.

10.1 (combined), against 1.97, 1.16, 1.01, and 2.18 for decentralized MPC. The ratios depend on the protocol and are not universal factors.

Mechanism of the noise sensitivity. Four factors contribute. (i) The output-bias correction of Eq. (31) uses the current measurement without filtering, so each noise sample shifts the entire predicted trajectory and is treated by the optimizer as a persistent disturbance. (ii) The tuned settings use $\lambda_u = 0.005$, which, by Eq. (32), yields a large noise-to-input gain; for MIMO-MPC this gain is a full matrix, so noise in one zone also moves neighboring heaters, which act back on the noisy zone through the plant coupling. Consistently, with $\lambda_u = 0.05$ in Study D the ratio of noisy to nominal RMSE is only 1.07 for MIMO-MPC and 1.02 for decentralized MPC, which supports, but does not prove, a tuning-related rather than an architectural cause. (iii) The nominal MIMO-MPC error is close to zero, so the same absolute noise-induced error appears as a much larger relative degradation. (iv) If the nominal case is evaluated with simulator-state prediction and the noisy cases with the measurement-based correction, part of the change in the ratio reflects this change of information setting.Filtering the bias, $\hat{\mathbf{d}}(k) = \alpha\hat{\mathbf{d}}(k-1) + (1-\alpha)[\mathbf{y}_{\mathrm{m}}(k) - \hat{\mathbf{y}}(k)]$, reduces the variance of the injected white noise by the factor $(1-\alpha)/(1+\alpha)$, for example to 5.3% for $\alpha = 0.9$, while preserving offset-free tracking of constant disturbances because the filter has unit static gain; the cost is a slower response to load changes, with an estimator time constant of $-T_s/\ln\alpha \approx 9.5$ s for $\alpha = 0.9$. A Kalman filter for a disturbance-augmented model [32, 38] and noise-aware tuning of λ_u are the systematic remedies; their evaluation is left for future work.

5.4 Scalability and Computational Cost

Across 9 to 36 zones (Fig. 7) the four-controller ordering is maintained, and MIMO-MPC reaches an RMSE of about 0.046–0.054 °C, with a QP success rate of 100% and no missed sampling deadline. In the nine-zone case, its RMSE is 1.9 times lower than

Table 6 Study C: tracking RMSE [°C] over 30 paired trials (mean ± standard deviation where stochastic); tuned settings of Study B

Scenario	Dec. MPC	MIMO-MPC	Ratio
S1 Nominal	0.0154	0.0024	6.32
S2 Noise	0.0302 ± 0.0011	0.0215 ± 0.0004	1.41
S3 Load	0.0179	0.0090	1.99
S4 Mismatch	0.0155 ± 0.0002	0.0028 ± 0.0001	5.62
S5 Combined	0.0335 ± 0.0017	0.0246 ± 0.0013	1.36

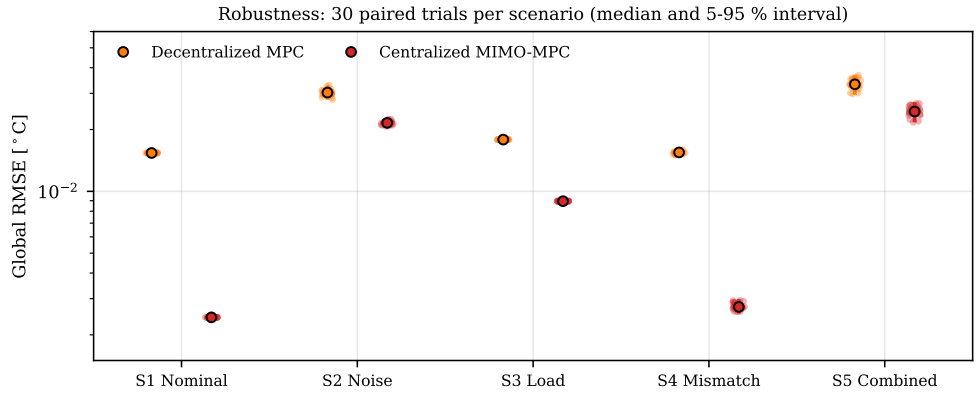

Fig. 6 Study C: paired-trial RMSE under nominal conditions, noise, load, mismatch, and their combination.

that of decentralized MPC under the combined perturbations. Because plant parameters, coupling rules, and reference amplitudes change with size, only comparisons between controllers within one size are meaningful.

Computation grows with dimension (Fig. 7b). The summed mean QP time of the n sequential decentralized problems rises from 6.8 to 25.1 ms and that of the single centralized QP from 2.1 to 21.3 ms, with log–log slopes of 0.95 and 1.67 with respect to N. The centralized controller stays within $T_s = 2$ s but grows faster, and its 95th-percentile time reaches 41.4 ms at 36 zones, against 26.4 ms for the sequential decentralized problems. Extrapolating the slopes places the crossover of the mean times near 45 zones, unless the sparse structure of the coupled QP is exploited by structure-exploiting or embedded solvers [39, 40]. The decentralized times are sequential sums; because the N local problems are independent, they could be solved in parallel, in which case the wall-clock time per step would be bounded below by the slowest local problem, of the order of the mean local time $25.1/36 \approx 0.7$ ms at 36 zones, excluding communication overhead. Parallel execution was not measured. These are desktop MATLAB results, not real-time guarantees. Cost also depends on tuning: for the tuned nine-zone controllers of Study B, the mean, 95th-percentile, and maximum times per step are 58.7, 68.3, and 96.7 ms for MIMO-MPC (270 decision variables) and

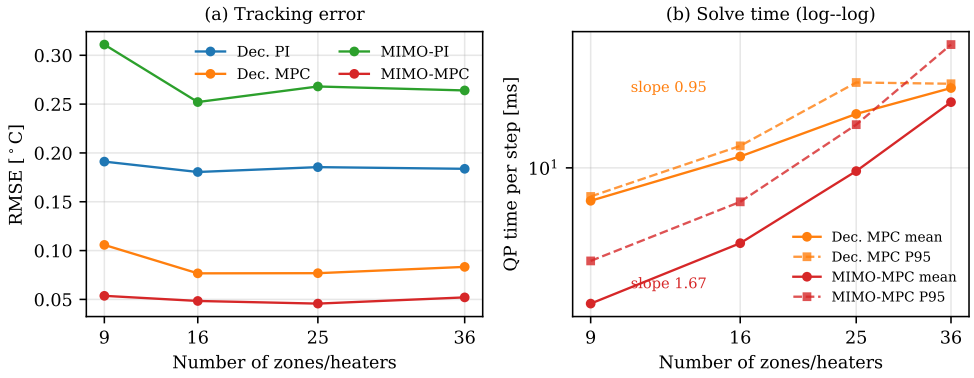

Fig. 7 Study D: (a) RMSE by controller and array size; (b) QP solution time per control step for decentralized (sequential sum) and centralized MPC. Larger grids are synthetic extensions of the nine-zone benchmark with edge-adjacent coupling only.

6.2, 6.7, and 22.8 ms for the nine sequential decentralized problems (9×10 variables), both with 100% success; the worst step uses 9.7% of $T_s = 1$ s. The PI controllers require 0.02–0.09 ms per step. The difference arises partly because the control horizon of the tuned MIMO-MPC is three times longer ($N_c = 30$ vs. 10); cost should therefore be reported with decision-vector size and tuning.

5.5 Model Consistency and Limitations

Re-simulating the logged inputs reproduces the recorded outputs, and the zero-coupling test agrees between centralized and decentralized prediction; these checks validate numerical consistency, not the physical fidelity of the model. The centralized controller uses the exact nominal coupled model whereas decentralized MPC keeps only direct-path dynamics, so the nominal results describe a favorable model-information setting for MIMO-MPC, not superiority under arbitrary uncertainty.

The main limitations are as follows. (1) *Plant fidelity.* The plants are synthetic LTI models. Radiation nonlinearity, operating-point dependence (Section 2.5), identification from measured TVAC data, and a validated state estimator are outside the benchmark, and the sensitivity of the conclusions to the specific cross-path rules (κ_τ, δ_d, γ_D/γ_N) has not been evaluated beyond the coupling sweep and the parametric mismatch scenarios. (2) *Attribution.* The matched sweep isolates the effect of coupling information as a whole, but static and dynamic coupling information have not been separated within the same MPC architecture; an MPC whose model contains the correct steady-state cross-gains but direct-path dynamics would be required. Moreover, MPC has future-reference information unavailable to PI, and PI is rule-tuned, so the four-controller comparison does not isolate coupling. (3) *References.* All references move the zones in the same direction, at moderate rates. Asynchronous, opposing, or fast references, for which interaction opposes rather than assists decentralized loops, have not been evaluated; given that the advantage of MIMO-MPC is largest for the

large, fast ramps of Study A and smallest for the small-target sweep, it is hypothesized that the benefit of coupling-aware prediction increases with reference rate, amplitude, and spatial heterogeneity. (4) *Noise and tuning.* Tuning uses noise-free trajectories, and the output correction is an unfiltered constant-bias disturbance model [33], so noise robustness may improve with other weights, bias filtering, or a Kalman filter. (5) *Output constraints.* The QP constrains actuator commands only. Temperature limits and inter-zone gradient limits, which are essential in TVAC test procedures, are not enforced predictively; the present controllers are applicable only where these limits are guaranteed by a supervisory interlock or by sufficiently conservative references. They can be added as soft constraints without losing convexity or feasibility (Section 3.2.5). (6) *Scalability.* The larger arrays use a different coupling rule and level than the nine-zone plant, so the effect of dimension is not isolated, and timing comes from a general-purpose MATLAB QP solver with sequential execution of the local problems; sparse optimization, parallel execution, and hardware-in-the-loop tests are left for future work.

Overall, the evidence for coupling-aware centralized MPC rests on the matched-setting comparisons and on the growth of its advantage with coupling strength. It persists under the tested disturbances and array sizes, but its magnitude depends on tuning, reference, and noise, and the absolute errors characterize an idealized benchmark.

6 Conclusion

This paper investigated coupling-aware centralized model predictive control for coordinated temperature regulation in a synthetic multi-zone TVAC benchmark. A heterogeneous MIMO first-order-plus-dead-time model with documented, physically motivated parameter rules was used to represent dynamic thermal interactions between neighboring heater–zone pairs. The control architecture incorporated these interactions into constrained multi-step prediction and was evaluated against decentralized PI, decentralized MPC, and static-decoupled MIMO-PI.

In the nominal nine-zone comparison, centralized MIMO-MPC achieved a tracking RMSE of 0.0008 °C, compared with 0.0097 °C for decentralized MPC, while satisfying the prescribed actuator constraints. This gap combines coupling information with the advantage of simulator-state prediction, and the millikelvin-level error is the floor of an idealized benchmark rather than achievable hardware accuracy. With a direct-path model, the joint QP decomposes exactly into the local problems, so under matched information and settings the advantage of the centralized controller can stem only from the coupling terms of its model. Coupling-sensitivity experiments provided further evidence of the role of interaction modeling: the two MPC architectures exhibited identical performance when cross-zone coupling was removed, whereas the advantage of centralized prediction increased to 12.8% with coupling strength under matched controller settings. Additional tests showed that the performance advantage persisted under model uncertainty and across synthetic arrays of up to 36 zones, although measurement noise reduced the relative improvement, mainly through the unfiltered

output-bias correction combined with an aggressive move penalty, and centralized optimization incurred increasing computational cost.

These findings indicate that explicitly accounting for dynamic cross-zone interactions can improve coordinated temperature tracking, but the magnitude of the benefit depends on coupling strength, controller tuning, reference trajectories, and operating conditions. The results are limited to physics-inspired synthetic models, with an exact nominal plant representation and desktop-based optimization. Future work will separate static and dynamic coupling information within the same MPC architecture, evaluate asynchronous and opposing references, filtered or Kalman-based disturbance estimation, and soft temperature and gradient constraints, and address experimentally identified thermal models, nonlinear heat-transfer effects, and hardware validation to establish the practical applicability of the approach.

References

[1] Almeida, J.S., Santos, M.B., Panissi, D.L., Garcia, E.C.: Effectiveness of low-cost thermal vacuum tests of a micro-satellite. Acta Astronautica **59**(6), 483–489 (2006) https://doi.org/10.1016/j.actaastro.2006.03.007

[2] Chisabas, R.S.S., Loureiro, G., Lino, C.d.O., Zabala, J.P.O., Salamanca, D.F.C.: Development of a thermal-vacuum chamber for testing in small satellites. In: 47th International Conference on Environmental Systems (2017). International Conference on Environmental Systems

[3] Salleh, N., Daud, S.M., Sabri, S.F., Ahmad-Salleh, N.A., Shariff, S.A., Md Yusof, Y., Adam, M.Z.: Enhancing temperature control method of thermal vacuum chamber for satellite testing using optimization algorithm: A review. Jurnal Teknologi **78**(5–7), 1–6 (2016) https://doi.org/10.11113/jt.v78.8820

[4] Salleh, N., Daud, S.M., Sabri, S.F., Abd Aziz, N.S.N.: Temperature control method optimization with global criterion genetic algorithm in thermal vacuum chamber for satellite testing. International Journal of Advanced and Applied Sciences **4**(12), 117–124 (2017) https://doi.org/10.21833/ijaas.2017.012.018

[5] Guo, J.-C., Li, F.-Y., Chen, A.-R., Zhang, L.-H., Liu, S.-W.: Automatic temperature control design for thermal vacuum tests based on fuzzy pid control. In: 2019 International Conference on Quality, Reliability, Risk, Maintenance, and Safety Engineering (QR2MSE), pp. 186–194 (2019). https://doi.org/10.1109/QR2MSE46217.2019.9021163 . IEEE

[6] Zhan, H., Sun, Y., Liu, D., Liu, H.: Adaptive neural network PID controller design for temperature control in vacuum thermal tests. In: 2016 28th Chinese Control and Decision Conference (CCDC), pp. 458–463 (2016). https://doi.org/10.1109/CCDC.2016.7531024 . IEEE

[7] Wu, D., Zhu, X., Wen, J., Zhu, L., Feng, Y.: Neural network based identification and self-tuning PID control for spacecraft thermal vacuum tests. In: Journal of

Physics: Conference Series, vol. 2216, p. 012058. IOP Publishing, ??? (2022). https://doi.org/10.1088/1742-6596/2216/1/012058

[8] Li, F., Yang, L., Ye, A., Zhao, Z., Shen, B.: Grouping neural network-based Smith PID temperature controller for multi-channel interaction system. Electronics **13**(4), 697 (2024) https://doi.org/10.3390/electronics13040697

[9] Chang, X., Zang, B., Gao, P., Li, X., Zhao, Z., Ye, A.: An average temperature feedback based PID controller for multi-channel balanced heating in thermal vacuum test. Journal of Computers **35**(5), 149–167 (2024) https://doi.org/10.53106/199115992024103505012

[10] Li, X., Chen, C.: Temperature prediction for aerospace thermal tests based on physical and LSTM hybrid model. Aerospace **11**(12), 964 (2024) https://doi.org/10.3390/aerospace11120964

[11] Zhang, J., Lu, T., Ye, T., Li, F., Zhang, Y., Liu, X.: Development of an intelligent control system for spacecraft vacuum thermal test. In: Journal of Physics: Conference Series, vol. 1848, p. 012144. IOP Publishing, ??? (2021). https://doi.org/10.1088/1742-6596/1848/1/012144

[12] Ding, W.J., Shan, W., Wang, Z.J., Liu, M.: Numerical model analysis and experimental verification of a thermal regulation strategy for thermal vacuum test. In: WIT Transactions on Modelling and Simulation, vol. 59, pp. 403–409. WIT Press, ??? (2015). https://doi.org/10.2495/CMEM150361

[13] Camacho, E.F., Bordons, C.: Model Predictive Control, 2nd edn. Advanced Textbooks in Control and Signal Processing. Springer, London (2004). https://doi.org/10.1007/978-0-85729-398-5

[14] Maciejowski, J.M.: Predictive Control with Constraints. Prentice Hall, Harlow, England (2002)

[15] Rawlings, J.B., Mayne, D.Q., Diehl, M.M.: Model Predictive Control: Theory, Computation, and Design, 2nd edn. Nob Hill Publishing, Madison, WI (2017)

[16] Qin, S.J., Badgwell, T.A.: A survey of industrial model predictive control technology. Control Engineering Practice **11**(7), 733–764 (2003) https://doi.org/10.1016/S0967-0661(02)00186-7

[17] Park, S.-w., Kim, S.: Nonlinear model predictive control of thermal vacuum chamber temperature. International Journal of Aeronautical and Space Sciences **25**(1), 213–228 (2024) https://doi.org/10.1007/s42405-023-00639-8

[18] Klupar, M.C., Larsson, D.T., Douglas, E.S.: Nonlinear Thermal Modeling and Predictive Control for Spacecraft Optical Temperature Regulation. Preprint (2026)

[19] Moroşan, P.-D., Bourdais, R., Dumur, D., Buisson, J.: Distributed MPC for multi-zone temperature regulation with coupled constraints. In: Proceedings of the 18th IFAC World Congress, vol. 44, pp. 1552–1557 (2011). https://doi.org/10.3182/20110828-6-IT-1002.02637 . IFAC

[20] Li, Z., Zhang, J.: Study on the distributed model predictive control for multi-zone buildings in personalized heating. Energy and Buildings **231**, 110627 (2021) https://doi.org/10.1016/j.enbuild.2020.110627

[21] Mork, M., Xhonneux, A., Müller, D.: Nonlinear distributed model predictive control for multi-zone building energy systems. Energy and Buildings **264**, 112066 (2022) https://doi.org/10.1016/j.enbuild.2022.112066

[22] Scattolini, R.: Architectures for distributed and hierarchical model predictive control – a review. Journal of Process Control **19**(5), 723–731 (2009) https://doi.org/10.1016/j.jprocont.2009.02.003

[23] Christofides, P.D., Scattolini, R., Peña, D., Liu, J.: Distributed model predictive control: A tutorial review and future research directions. Computers & Chemical Engineering **51**, 21–41 (2013) https://doi.org/10.1016/j.compchemeng.2012.05.011

[24] Pedersen, T.H., Hedegaard, R.E., Knudsen, M.D., Petersen, S.: Comparison of centralized and decentralized model predictive control in a building retrofit scenario. In: Energy Procedia, vol. 122, pp. 979–984. Elsevier, ??? (2017). https://doi.org/10.1016/j.egypro.2017.07.411

[25] Zhao, J., Xu, D., Yuan, X., Liu, D.: Cooling supply allocation model predictive control strategy of data center computer room air conditioning systems considering multi-zone thermal coupling effects. Applied Thermal Engineering **277**, 127103 (2025) https://doi.org/10.1016/j.applthermaleng.2025.127103

[26] Mohammadi, A., Kramer, R., Nagy, Z.: Coordination architecture shapes continuous demand response outcomes in building districts. In: Proceedings of the 13th ACM International Conference on Systems for Energy-Efficient Buildings, Cities, and Transportation (BuildSys '26). ACM, Banff, AB, Canada (2026). https://doi.org/10.1145/3744256.3812578

[27] Michailidis, P., Michailidis, I., Minelli, F., Çoban, H.H., Kosmatopoulos, E.: Model predictive control for smart buildings: Applications and innovations in energy management. Buildings **15**(18), 3298 (2025) https://doi.org/10.3390/buildings15183298

[28] Skogestad, S., Postlethwaite, I.: Multivariable Feedback Control: Analysis and Design, 2nd edn. John Wiley & Sons, Chichester, UK (2005)

[29] Skogestad, S.: Simple analytic rules for model reduction and PID controller tuning. Journal of Process Control **13**(4), 291–309 (2003) https://doi.org/10.1016/S0959-1524(02)00062-8

[30] Incropera, F.P., DeWitt, D.P., Bergman, T.L., Lavine, A.S.: Fundamentals of Heat and Mass Transfer, 6th edn. John Wiley & Sons, Hoboken, NJ (2007)

[31] Gilmore, D.G. (ed.): Spacecraft Thermal Control Handbook, Volume I: Fundamental Technologies, 2nd edn. The Aerospace Press, El Segundo, CA (2002)

[32] Muske, K.R., Badgwell, T.A.: Disturbance modeling for offset-free linear model predictive control. Journal of Process Control **12**(5), 617–632 (2002) https://doi.org/10.1016/S0959-1524(01)00051-8

[33] Pannocchia, G., Rawlings, J.B.: Disturbance models for offset-free model-predictive control. AIChE Journal **49**(2), 426–437 (2003) https://doi.org/10.1002/aic.690490213

[34] Maeder, U., Borrelli, F., Morari, M.: Linear offset-free model predictive control. Automatica **45**(10), 2214–2222 (2009) https://doi.org/10.1016/j.automatica.2009.06.005

[35] Scokaert, P.O.M., Rawlings, J.B.: Feasibility issues in linear model predictive control. AIChE Journal **45**(8), 1649–1659 (1999) https://doi.org/10.1002/aic.690450805

[36] Kerrigan, E.C., Maciejowski, J.M.: Soft constraints and exact penalty functions in model predictive control. In: Proceedings of the UKACC International Conference (Control 2000), Cambridge, UK (2000)

[37] International Electrotechnical Commission: IEC 60751:2022 – Industrial platinum resistance thermometers and platinum temperature sensors. International Standard, Geneva (2022)

[38] Kuntz, S.J., Rawlings, J.B.: Maximum likelihood identification of linear models with integrating disturbances for offset-free control. IEEE Transactions on Automatic Control **70**(9), 5675–5689 (2025) https://doi.org/10.1109/TAC.2025.3547607

[39] Ferreau, H.J., Kirches, C., Potschka, A., Bock, H.G., Diehl, M.: qpOASES: a parametric active-set algorithm for quadratic programming. Mathematical Programming Computation **6**(4), 327–363 (2014) https://doi.org/10.1007/s12532-014-0071-1

[40] Stellato, B., Banjac, G., Goulart, P., Bemporad, A., Boyd, S.: OSQP: an operator splitting solver for quadratic programs. Mathematical Programming Computation **12**(4), 637–672 (2020) https://doi.org/10.1007/s12532-020-00179-2