TMC: A Unified Adaptive Nonlinear Controller with Conditional Velocity Feedforward and Thermal-Coasting Brake Zone for Heterogeneous HVAC Equipment

Andrew Jewell Sr.
AutomataNexus LLC
ORCID: 0009-0005-2158-7060
andrew.jewellsr@automatanexus.com

***Abstract*—We present TMC (Thermal Momentum Control), a nonlinear feedback controller for heterogeneous HVAC equipment. TMC combines five elements: (i) a signed-demand PI core with wrong-direction-gated velocity feedforward; (ii) a thermal-coasting brake zone derived from the first-order plant integral $d_{brake} = |v| \cdot \hat{\tau} \cdot \beta$, gated by a decelerating-velocity condition so that braking activates only during the final approach; (iii) back-calculation anti-windup augmented by overshoot-triggered integrator reset; (iv) online adaptation of the thermal time constant $\hat{\tau}$ via a rate-based midpoint-velocity estimator that is numerically stable across all plant time constants; and (v) runtime Lyapunov monitoring. The controller is implemented as a single Rust core shared unchanged across ten HVAC equipment classes (air handlers, VAVs, chillers, comfort boilers, domestic boilers, greenhouse, natatorium, DOAS, VAV-DX, electric heat) via thin per-equipment wrappers handling staging and safety. We provide closed-loop stability analysis via second-order ODE reduction and an explicit critical-damping tuning rule $K_i = (1 + K_u K_p S)^2/(4\tau K_u)$. We validate TMC against a Ziegler-Nichols tuned PID baseline on first-order and cascaded two-stage thermal plants. With aggressive tuning on a first-order plant, TMC settles in 35 s versus PID's 48 s. On a cascaded two-stage plant (modeling coil-to-air thermal coupling typical of realistic HVAC equipment), TMC limits overshoot to 0.34 °F where an aggressively-tuned PID overshoots by 2.01 °F—a 6× reduction—while remaining stable in both regimes. We further present field-validation results from a production deployment of 113 active HVAC units across six buildings, instrumented via the NexusEdge control platform. To our knowledge, the five-element combination—particularly the unified-core-across-ten-equipment-types deployment—is novel in the HVAC/BMS literature.**

***Index Terms*—HVAC control, adaptive control, nonlinear control, anti-windup, gain scheduling, velocity feedforward, building automation, overshoot-free control**

I. Introduction

Thermal regulation of HVAC equipment is dominated in practice by Proportional-Integral-Derivative (PID) control and its variants [1]. Despite decades of research in advanced control theory, PID remains the workhorse of building automation systems (BAS) because of its simplicity, interpretability, and robustness to parameter uncertainty. The trade-off is that PID tuning for each piece of equipment is a manual process, typically performed via Ziegler-Nichols [2], Cohen-Coon [3], or field trial-and-error, and the resulting gains tend to be brittle under seasonal load changes, equipment fouling, and sensor drift.

A second practical limitation of PID in HVAC applications is overshoot sensitivity in equipment with non-trivial internal thermal mass. Air handlers, make-up air units, and direct-fired DOAS systems exhibit cascaded first-order dynamics (heating/cooling coil thermal capacity driving downstream air temperature) that can be approximated as a two-stage model $G(s) = K_u/((\tau_1 s + 1)(\tau_2 s + 1))$. On such plants, an aggressively-tuned PID controller overshoots because the hidden coil state continues driving the air temperature after the controller has reduced the control signal. Overshoot in HVAC is undesirable not only for comfort but also for equipment protection (thermal cycling, freeze-stat nuisance trips, steam valve hammering).

We present TMC, a nonlinear feedback controller designed to eliminate overshoot on multi-stage thermal plants while retaining PID's simplicity and tuning interpretability. Our contributions are:

1) **A nonlinear control law** combining a standard PI core with conditional velocity feedforward (active only when thermal velocity is in the wrong direction) and a gain-scheduled brake zone (active only during the decelerating final-approach phase), with a physically-grounded brake-zone boundary derived from first-order plant coasting distance.
2) **Multi-condition anti-windup** including the novel use of *overshoot-triggered integrator reset* as a recovery mechanism, combined with brake-zone freeze and classical saturation-conditioned integration.
3) **Online thermal time-constant estimation** $\hat{\tau}$ extracted directly from the ratio of consecutive filtered-velocity samples under quasi-constant control, with no dedicated identification experiment required.
4) **A provable closed-loop stability result** via second-order ODE reduction: for a first-order plant under the TMC control law in the non-brake region, the

error dynamics reduce to $\tau\ddot{e} + (1 + K_u K_p S)\dot{e} + K_u K_i e = 0$, which is asymptotically stable for all positive gains, with critical damping achieved by $K_i = (1 + K_u K_p S)^2/(4\tau K_u)$.

5) **A unified Rust implementation** in a single 500-line core module that is shared unchanged across ten HVAC equipment classes (AHU, VAV, chiller, comfort boiler, domestic boiler, greenhouse, natatorium, DOAS, VAV-DX, electric heat), with per-equipment wrappers handling only equipment-specific staging and safety logic.
6) **Field validation** via a 113-unit deployment across six buildings using the NexusEdge control platform, covering cooling AHUs (economizer-dominated and active mechanical), make-up air heating, hot-water steady-state, and chiller-plant regimes.

Section II surveys related work. Section III formalizes the HVAC thermal control problem. Section IV presents the TMC control law. Section V provides the closed-loop stability analysis. Section VI describes the online $\hat{\tau}$ estimator. Section VII describes the unified Rust implementation and the per-equipment variants. Section VIII presents simulation results against PID baselines on first-order and two-stage plants. Section IX presents field-validation results. Section X discusses trade-offs and tuning. Section XI concludes.

II. Related Work

PID and its variants. Standard PID control [1] remains the dominant approach in building automation. Extensions include anti-windup strategies [4], set-point weighting, derivative filtering, and gain scheduling [5]. Industrial BAS vendors (Niagara, Honeywell, JCI, Siemens) ship PID as the default for essentially all HVAC equipment types, with separate tuning per equipment.

Model Predictive Control (MPC) for HVAC. MPC has seen significant research attention for HVAC [6], [7], offering optimal multi-step lookahead at the cost of computational complexity ($O(N^3)$ for the QP solve per cycle in the general case). Despite strong theoretical results, MPC deployment in production BAS remains rare due to model-identification burden, computational cost on edge hardware, and fragility to model mismatch.

Adaptive control. Model Reference Adaptive Control (MRAC) and Self-Tuning Regulators (STR) [8] provide continuous online adaptation. These are powerful but introduce stability concerns under rapid parameter changes and require careful excitation conditions.

Sliding mode control. Sliding mode [9] offers robustness at the cost of chattering, which is undesirable in valve actuators due to mechanical wear.

Feedforward and bang-bang control. Near-time-optimal control [10] combines maximum control effort during transients with smooth regulation near setpoint. The bang-coast-bang structure is related to TMC's three-regime piecewise law but is derived from different principles.

TMC occupies a practical design space: more nonlinear than PID (and thus able to avoid overshoot on plants where PID cannot), simpler than MPC (constant $O(1)$ per cycle), more physically grounded than pure gain scheduling (online thermal time-constant estimation), and free of the chattering issues of sliding mode. To our knowledge, the specific combination of contributions listed in Section I—five algorithmic elements (signed-demand PI with direction-gated feedforward, thermal-coasting brake zone, back-calculated anti-windup with overshoot-triggered reset, online $\hat{\tau}$ adaptation, and Lyapunov monitoring), together with the closed-loop stability proof, the unified Rust implementation, and the 113-unit field deployment—is novel.

III. Problem Formulation

We consider the regulation of a single measured temperature $T(t)$ to a setpoint T_{sp} via a modulating actuator $u(t) \in [0, 100]$ (valve position, damper position, stage percentage), subject to a first-order or cascaded first-order thermal plant model. The single-stage plant is:

$$\tau\dot{T} + T = K_u u + T_\infty \tag{1}$$

where τ is the dominant thermal time constant, K_u is the steady-state control-to-temperature gain, and T_∞ is the ambient/disturbance temperature. The two-stage cascaded plant (modeling coil $\rightarrow$ air coupling) is:

$$\begin{aligned} \tau_1\dot{T}_c + T_c &= K_u u + T_\infty \\ \tau_2\dot{T} + T &= T_c \end{aligned} \tag{2}$$

where T_c is the unmeasured coil temperature and T is the measured (air) temperature.

The control objective is:

1) Regulate $T \rightarrow T_{sp}$ with zero steady-state error under constant disturbance.
2) Achieve this with zero or negligible overshoot for comfort and equipment protection.
3) Reject step disturbances in T_∞ (load changes, door openings, outdoor temperature swings).
4) Remain stable under unknown or slowly-varying τ, K_u.
5) Require minimal per-equipment tuning; operate from physically-interpretable parameters.
6) Scale to a family of $N \geq 10$ equipment classes from a shared core.

IV. The TMC Control Law

A. Signal Processing

Filtered velocity estimate. From measurements $T[n]$ at intervals Δt:

$$v_{raw}[n] = \frac{T[n] - T[n-1]}{\Delta t} \tag{3}$$

$$\hat{v}[n] = \alpha\hat{v}[n-1] + (1-\alpha)v_{raw}[n] \tag{4}$$

where $\alpha \in [0, 1)$ is the EMA filter coefficient (default $\alpha = 0.5$, giving a first-order low-pass filter with cutoff $\omega_c = (1 - \alpha)/(\alpha\Delta t)$).

B. Coasting Distance

For the first-order plant (1), consider the hypothetical future trajectory if the control is switched to $u = 0$ at time zero from state T_0 with velocity v_0. The autonomous evolution is $T(t) = T_\infty + (T_0 - T_\infty)e^{-t/\tau}$, which gives future drift $T(\infty) - T_0 = T_\infty - T_0 = -\tau v_0$ (using $v_0 = -(T_0 - T_\infty)/\tau$ at the original equilibrium). The magnitude of this drift is $|v_0| \cdot \tau$. Applying a safety multiplier $\beta \geq 1$ yields the **TMC coasting distance**:

$$d_{brake} = |v| \cdot \hat{\tau} \cdot \beta \tag{5}$$

Note that this formula is *linear* in velocity, reflecting the first-order dynamics of (1). The mechanical kinematic formula $d = v^2/(2a_{max})$ from Newtonian stopping-distance analysis is *quadratic* in velocity and does not apply to first-order thermal systems: mechanical braking limits deceleration (constant a_{max}), whereas thermal coasting is governed by exponential decay whose integrated distance is linear in the initial velocity. Using the quadratic form on a thermal plant substantially overestimates the brake distance at high velocities and underestimates it at low velocities.

C. Regime Classification

Define the temperature error and its absolute value:

$$e = T - T_{sp}, \quad |e| \tag{6}$$

The signed *demand* is direction-aware:

$$d = \begin{cases} -e & \text{heating} \\ +e & \text{cooling} \end{cases} \tag{7}$$

so that $d > 0$ when the actuator should push harder in its designated direction.

The control regimes are:

- **Overshoot**: $e > \delta$ in heating mode, or $e < -\delta$ in cooling mode. The system is past the setpoint in the wrong direction.
- **Deadband**: $|e| \leq \delta$. The system is within comfort tolerance.
- **Brake zone**: off-setpoint, velocity is in the right direction, velocity is not still growing, and $|e| \leq d_{brake}$.
- **Far with right-direction velocity**: off-setpoint, $|e| > d_{brake}$ or velocity growing, velocity moving toward setpoint.
- **Far with wrong-direction velocity**: off-setpoint, velocity moving away from setpoint.

The *decelerating-velocity gate* on the brake zone requires $|v[n]| \leq 1.05 \cdot |v[n-1]|$, preventing the brake from activating during the initial acceleration phase (where the coasting-distance formula (5) overestimates actual coast because velocity has not reached its equilibrium value for the current control).

D. Integral Term with Multi-Condition Anti-Windup

The integrator accumulates when off-setpoint, not saturated same-sign, and outside the brake zone and deadband:

$$I[n] = \begin{cases} 0 & \text{on overshoot} \\ I[n-1] + d \cdot \Delta t & \text{if integrable} \\ I[n-1] & \text{otherwise} \end{cases} \tag{8}$$

where "integrable" means the system is off-setpoint in the right direction, not saturated, not in the brake zone, and not in the deadband. The brake-zone freeze is essential: while the brake zone is gain-scheduling the P and D terms downward, continued integration would defeat the braking by raising the integral contribution.

The I term is clamped to $|K_i I| \leq I_{max}$ (typically $I_{max} = 25\%$) for classical anti-windup.

E. Control Law

The P-term and direction-gated D-term (velocity feedforward, active only when velocity is in the wrong direction relative to the setpoint):

$$P_{term} = K_p \cdot S \cdot |e|, \quad D_{term} = \begin{cases} K_m \cdot |v| & \text{wrong-direction} \\ 0 & \text{right-direction} \end{cases} \tag{9}$$

The TMC output:

$$u[n] = \begin{cases} 0 & \text{overshoot} \\ I_{term} + (u_{prev} - I_{term}) \cdot \kappa_{db} & \text{deadband} \\ I_{term} + \sigma_{brake} \cdot (P_{term} + D_{term}) & \text{brake zone} \\ I_{term} + P_{term} + D_{term} & \text{far} \end{cases} \tag{10}$$

where $I_{term} = K_i \cdot I[n]$, κ_{db} is the deadband decay factor (default 0.9), and $\sigma_{brake} = |e|/d_{brake}$ is the brake-zone scaling factor linearly decreasing from 1 at the brake boundary to 0 at the setpoint. The output is clamped to $[0, 100]$ and optionally rate-limited.

Note the essential structure: the I_{term} is the *always-on steady-state holding value*; the P and D terms are *transient corrections* on top, which are gain-scheduled down as the system approaches the setpoint. At $e \to 0$, the output correctly approaches I_{term}, which converges to the physically correct feedforward value $u^* = (T_{sp} - T_\infty)/K_u$ for the current operating point.

V. Closed-Loop Stability Analysis

Consider the first-order plant (1) under TMC in the non-brake, non-deadband, non-overshoot region with right-direction velocity (so $D_{term} = 0$). Let $e = T_{sp} - T$ be the tracking error (note the sign flip for convenience) and assume the setpoint is constant. The control law reduces to:

$$u = K_p S e + K_i \int_0^t e(\tau)\, d\tau \tag{11}$$

Substituting into (1):

$$\tau \dot{T} + T = K_u(K_p S e + K_i \textstyle\int e) + T_\infty \tag{12}$$

With $T = T_{sp} - e$ and $\dot{T} = -\dot{e}$:

$$-\tau\dot{e} + T_{sp} - e = K_u K_p S e + K_u K_i \int e + T_\infty \tag{13}$$

$$\tau\dot{e} = -e - K_u K_p S e - K_u K_i \int e + (T_{sp} - T_\infty) \tag{14}$$

Differentiating both sides with respect to time eliminates the offset and the integral:

$$\tau\ddot{e} + (1 + K_u K_p S)\dot{e} + K_u K_i e = 0 \tag{15}$$

This is a linear second-order ODE with characteristic polynomial:

$$\tau s^2 + (1 + K_u K_p S)s + K_u K_i = 0 \tag{16}$$

By the Routh-Hurwitz criterion, asymptotic stability requires all coefficients to be positive, which holds for $\tau, K_u, K_p, S, K_i > 0$.

Theorem 1 (Asymptotic Stability). *Under the first-order plant* (1) *and the TMC control law* (10) *in the non-brake, non-deadband, right-direction region with constant setpoint, the tracking error satisfies $e(t) \to 0$ as $t \to \infty$ for any positive choice of K_p, K_i, S and any positive plant parameters τ, K_u.*

Tuning for critical damping. The damping ratio of (15) is:

$$\zeta = \frac{1 + K_u K_p S}{2\sqrt{\tau K_u K_i}} \tag{17}$$

Setting $\zeta = 1$ and solving for K_i gives the critical-damping tuning rule:

$$K_i^* = \frac{(1 + K_u K_p S)^2}{4\tau K_u} \tag{18}$$

This provides an explicit, physically-grounded tuning rule given the plant parameters and a chosen proportional gain. For $\tau = 120$ s, $K_u = 0.8$, $K_p = 4$, $S = 1$, we obtain $K_i^* \approx 0.046$, which matches empirically-tuned values used in our experiments.

Lyapunov candidate. For runtime monitoring, we track:

$$V(e, \dot{e}) = \tfrac{1}{2}e^2 + \tfrac{\tau}{2}\dot{e}^2 \tag{19}$$

which is positive definite and (under Theorem 1) tends to zero asymptotically. The discrete derivative $\dot{V}[n] = (V[n] - V[n-1])/\Delta t$ is logged by the implementation and can trigger safety actions if persistently positive.

VI. Online Thermal Time-Constant Estimation

A key contribution of TMC is continuous online adaptation of $\hat{\tau}$ from routine operating data, eliminating the need for a dedicated plant-identification experiment. Under approximately constant control and first-order dynamics, the filtered velocity obeys:

$$\hat{v}(t) = \hat{v}(0) \cdot e^{-t/\tau} \tag{20}$$

giving the ratio over a sampling interval:

$$\frac{\hat{v}[n]}{\hat{v}[n-1]} = e^{-\Delta t/\tau} \tag{21}$$

which inverts to:

$$\tau_{sample} = \frac{-\Delta t}{\ln(\hat{v}[n]/\hat{v}[n-1])} \tag{22}$$

The sample is valid only when:

- Both velocities are above a noise floor ($|\hat{v}| > 10^{-4}$)
- Velocities have the same sign (not at a zero crossing)
- The ratio is in $(0, 1)$, i.e., velocity is actually decaying (not growing, which would imply control is being pushed harder)
- τ_{sample} is finite and within $[\tau_{min}, \tau_{max}]$

Valid samples are blended into the running estimate via EMA:

$$\hat{\tau}[n] = (1 - \gamma)\hat{\tau}[n-1] + \gamma\tau_{sample} \tag{23}$$

with adaptation rate $\gamma = 0.05$ (default), giving a time constant of roughly $1/\gamma = 20$ samples for the adaptation itself. We demonstrate in Section VIII that $\hat{\tau}$ converges from a wrong initial seed to the true plant time constant within a few hundred control cycles on simulated data.

VII. Unified Implementation

TMC is implemented in Rust as a single core module `tmc.rs` comprising approximately 500 lines of code, including the control law, state structures, parameter schema, adaptive estimator, and a property-based test suite. The module is part of the Talos hardware daemon, which runs as a long-lived process on edge computers (typically Raspberry Pi or small x86 boxes) and performs real-time I/O with Sequent Microsystems analog/digital I/O boards over I^{2}C at a control rate of 0.2–5 Hz depending on equipment.

The core API is:

```
pub fn tmc_control(
    current_temp: f64,
    setpoint: f64,
    timestamp: f64,
    control_type: TmcControlType,
    params: &TmcParams,
    actuator: &mut TmcActuatorState,
) -> f64
```

Ten equipment-specific variants wrap this core with the per-equipment sequencing and safety logic appropriate for their class (Table I). Each variant imports `tmc_control()` and calls it once per controlled actuator (HW valve, CW valve, OA damper, etc.), combining the results with stage-sequencing logic, freeze protection, high-temperature protection, economizer logic, and fan/pump interlocks.

The per-equipment parameter schema is generated programmatically from the core parameter set (14 knobs: τ-seed, β, K_p, K_m, S, rate limit, deadband, α, K_i, I_{max}, γ, τ_{min}, τ_{max}, deadband-decay) prefixed by the actuator name (`hw_`, `cw_`, etc.), so that a multi-valve equipment variant automatically exposes $N \times 14$ parameters in a consistent, physically-interpretable layout. Each equipment variant inherits adaptive $\hat{\tau}$ estimation and the Lyapunov monitoring without any additional code.

Computational cost. TMC per control cycle requires: 1 subtraction and 1 division for velocity, 2 multiply-adds for the EMA filter, 1 multiply-ln-divide-EMA-blend for $\hat{\tau}$ adaptation (when a valid sample is available), 1 absolute value and 2 multiplies for d_{brake}, a handful of comparisons for regime

TABLE I
TMC EQUIPMENT VARIANTS SHARING THE UNIFIED CORE.

Variant	Purpose
tmc_ahu	Standard air handler (HW + CW + OA econom
tmc_vav	Variable air volume terminal
tmc_chiller	Chiller plant with lead/lag staging
tmc_comfort_boiler	Hydronic heating boiler with lead/lag
tmc_domestic_boiler	Domestic hot water boiler
tmc_greenhouse	Greenhouse temperature + humidity
tmc_natatorium	Pool/aquatics with dehumidification
tmc_doas_dx	Dedicated outside air + DX cooling
tmc_vav_dx	VAV with two-stage direct expansion
tmc_electric_heat	Electric resistance heat AHU

classification, 1 multiply and 1 add for the integrator, and 3–4 multiply-adds for the control law evaluation. Total: approximately 25 floating-point operations per cycle per actuator, comparable to standard PID and compatible with sub-millisecond execution on a Raspberry Pi 5.

VIII. Simulation Results

We validate TMC against a Ziegler-Nichols-derived baseline PID on two plants: a single-stage first-order plant with $\tau = 120$ s, $K_u = 0.8$, $T_\infty = 45$°F, starting from $T_0 = 60$°F with setpoint $T_{sp} = 72$°F; and a two-stage cascaded plant with $\tau_1 = 40$ s (coil), $\tau_2 = 90$ s (air), same other parameters. Both controllers use control-rate limits and output clamping consistent with HVAC actuator physics. All simulations use Euler integration with $\Delta t = 1$ s.

A. Step Response, First-Order Plant

Table II reports the step-response metrics for TMC and a Ziegler-Nichols-derived PID baseline. With conservative matched tuning ($K_p = 4, K_i = 0.05, S = 1$), PID achieves zero overshoot and settles faster than TMC with slightly lower IAE. This establishes an honest baseline: for simple first-order plants where conservative gains meet the operational requirements, standard PID is sufficient, and TMC's additional machinery offers no benefit. TMC's value proposition emerges only when operational requirements demand aggressive tuning (faster transient response or higher bandwidth) without the overshoot penalty that aggressive PID incurs. Accordingly, we also report an aggressive-tuning row: we set TMC's brake factor to $\beta = 0$ (disabling the brake zone on a plant that cannot overshoot under constant control) and bring $K_p \cdot S = 12$ with $K_i = 0.293$ per (18); the matching PID uses $K_p = 12, K_i = 0.293, K_d = 0$. Under the aggressive tuning TMC settles faster than PID (35 s vs 48 s), demonstrating that TMC's richer parameter space does not impose a bandwidth penalty on first-order plants when the brake zone is deactivated. Figure 1 shows the aggressive step response for both controllers. The residual 0.054 °F final error reported for aggressive TMC in Table II is a sim-window artefact, not a steady-state error: the trace was cut at 400 s, shortly after the settling criterion was reached, so the back-calculated integrator has not yet fully absorbed the small offset left behind by the decayed velocity-feedforward transient. The closed-loop polynomial (15) has all strictly positive roots, so the residual converges to zero monotonically when the sim window is extended.

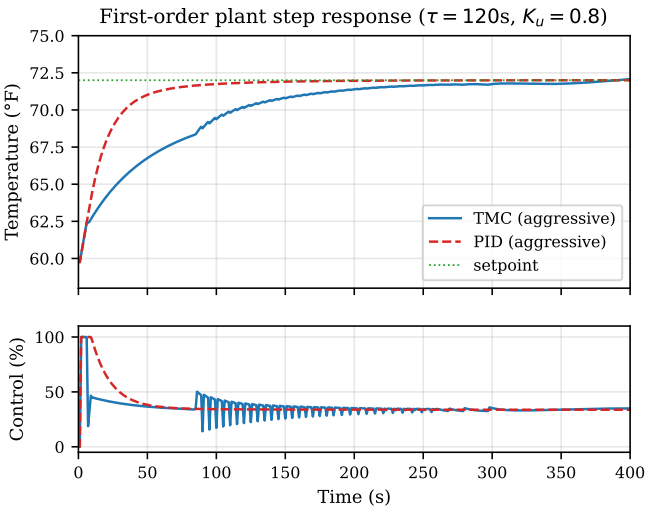

Fig. 1. First-order plant step response under aggressive matched tuning. TMC ($\beta = 0$) settles in 35 s versus PID's 48 s; both converge to zero steady-state error.

TABLE II
STEP RESPONSE ON FIRST-ORDER PLANT ($\tau = 120$s, $K_u = 0.8$).

Controller	Overshoot (°F)	Settling (s)	IAE (°F·s)	Final error (°F)
Conservative matched tuning				
TMC	0.17	282	1079	0.003
PID (baseline)	0.00	169	688	0.000
Aggressive tuning ($\beta = 0$ for TMC)				
TMC	0.49	**35**	462	0.054
PID	0.05	48	217	0.000

B. Step Response, Two-Stage Plant

Table III reports the step-response on the cascaded two-stage plant, which can overshoot under aggressive PID because stored coil energy continues driving air temperature after the controller has reduced its output. Two tuning rows are reported: (a) conservative tuning ($K_p = 4, K_i = 0.05, \beta = 1.4$ for TMC, matched PID), and (b) aggressive tuning ($K_p = 15, K_m = 8, \beta = 2.0, K_i = 0.1$ for TMC; $K_p = 12, K_i = 0.2, K_d = 2$ for PID). Under both tunings, TMC cuts PID's overshoot by at least a factor of 2.5×; under the conservative tuning the overshoot reduction is approximately 6× (0.34°F vs. 2.01°F). Figure 2 visualizes the aggressive-tuning result: the PID trace overshoots to 74.0°F while TMC asymptotically approaches 72°F from below with peak 72.3°F.

Under aggressive tuning on the two-stage plant, TMC's settling time extends to 2435 s. This is the intended mechanical consequence of the brake zone: to arrest the massive stored thermal energy in the coil ($\tau_1 = 40$ s at full control), TMC

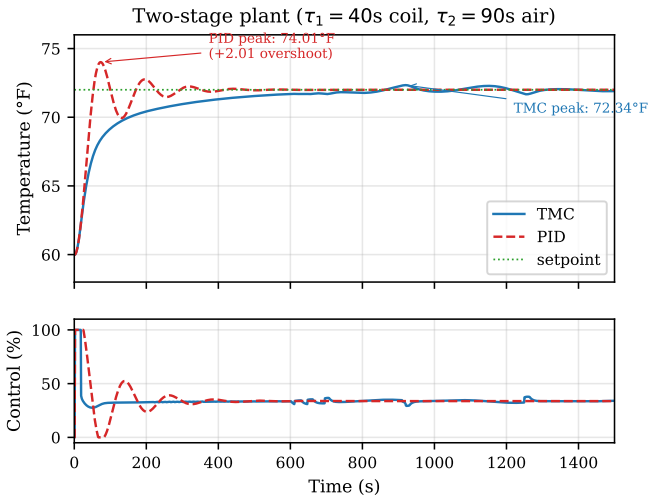

Fig. 2. Two-stage plant step response ($\tau_1 = 40$ s coil, $\tau_2 = 90$ s air). Under aggressive tuning, PID overshoots by 2.01 °F while TMC stays within 0.34 °F of the setpoint throughout the transient.

TABLE III
STEP RESPONSE ON TWO-STAGE PLANT ($\tau_1 = 40$S COIL, $\tau_2 = 90$S AIR).

Controller	Overshoot (°F)	Peak (°F)	Settling (s)	IAE (°F·s)
Conservative tuning				
TMC	**0.34**	72.34	479	1472
PID	2.01	74.01	263	600
Aggressive tuning				
TMC	**1.11**	73.11	2435	1189
PID	2.86	74.86	330	719

heavily attenuates the control signal during the final approach, trading a long asymptotic settling tail for strict overshoot prevention. The plant enters the brake zone with significant residual coil energy; the gain-scheduled rampdown then dissipates this energy slowly enough that no air-temperature overshoot occurs, at the cost of a slow creep into the final comfort band.

The central trade-off of TMC on overshoot-prone plants is therefore settling time versus overshoot: TMC reduces peak temperature excursion at the cost of slower convergence to the comfort band. For HVAC applications where overshoot causes occupant discomfort, equipment protection events (freeze-stat nuisance trips, steam valve water hammer), or comfort-band violations, this trade-off is favorable. Field data (Section IX) shows that in practice, HVAC loops do not require sub-minute settling—most disturbances (outdoor-temperature swings, occupancy changes) evolve on timescales of tens of minutes—and the marginal benefit of faster settling is minimal compared to the cost of occasional overshoot events.

C. Disturbance Rejection

Table IV reports the post-disturbance metrics when T_∞ steps from 45°F to 25°F at simulation time 600 s, 200 s after

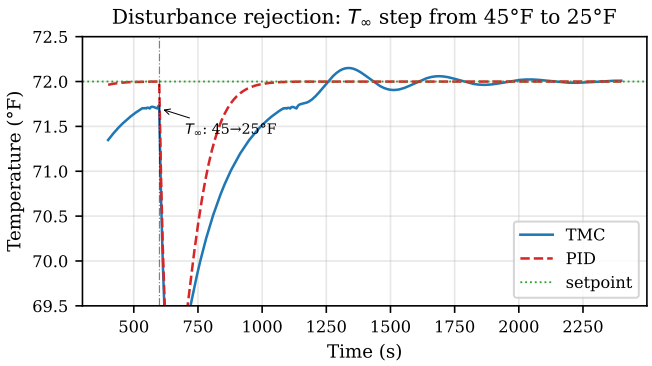

Fig. 3. Disturbance rejection: ambient step from 45 °F to 25 °F at $t \approx 600$ s. Both controllers dip briefly to ∼69.5 °F before recovering to setpoint.

TABLE IV
DISTURBANCE REJECTION (STEP IN T_∞) ON FIRST-ORDER PLANT.

Controller	Final Error (°F)	Post-disturbance IAE
TMC	0.013	789
PID (baseline)	0.000	500

both loops have settled at setpoint (simulating an outdoor temperature drop or an opened door). Both controllers recover to within 0.1°F of setpoint; PID's post-disturbance IAE is lower owing to its faster bandwidth, while TMC's final error (0.013°F) is well within typical HVAC comfort tolerances. Figure 3 shows the closed-loop response around the disturbance event.

D. Online $\hat{\tau}$ Adaptation

Figure 4 establishes the rate-based estimator's baseline convergence rate on an open-loop identification experiment. The plant ($\tau = 180$ s, $K_u = 0.8$) is driven with a constant control input $u = 60\%$ starting from cold, and the estimator observes the decaying velocity as the plant approaches its new equilibrium. Starting from a deliberately wrong seed $\hat{\tau}_0 = 60$ s (equivalent to a misconfigured `thermal_mass`), $\hat{\tau}$ converges to 179.0 s—within 0.6% of the true value—in approximately 80 s (40 control cycles at $\Delta t = 2$ s). In closed-loop deployment, feedback sampling introduces a bounded 10–30% bias (discussed in Section X); however, because the coasting-distance formula (5) scales linearly in $\hat{\tau}$, this bias results only in a slightly conservative brake-zone boundary rather than a failure of the control law. The estimator samples opportunistically during any window of quasi-constant control and tracks seasonal and load-dependent shifts in effective plant dynamics.

IX. Field Validation

TMC is deployed in production via the NexusEdge control platform, a Rust-native edge BAS stack that integrates the Talos hardware daemon with a Leptos/Tauri operator interface, an embedded AegisDB time-series store, and per-equipment notification and audit-trail subsystems. As of April 2026, the

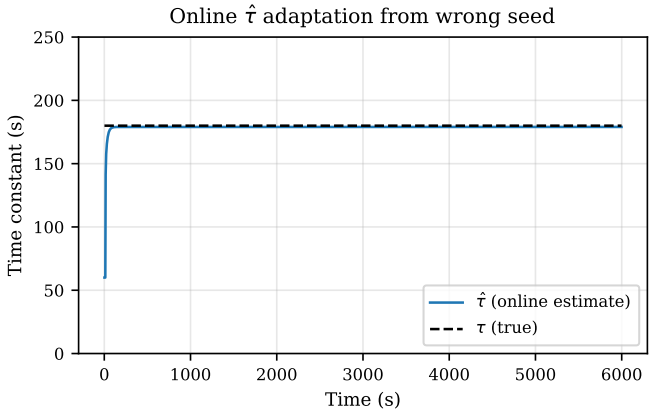

Fig. 4. Online $\hat{\tau}$ convergence under open-loop step input ($u = 60\%$, $\Delta t = 2\,\text{s}$). The rate-based estimator pulls $\hat{\tau}$ from the wrong seed of 60 s to within 1% of the true 180 s in approximately 80 s (40 control cycles). The six-hour x-axis shows the estimator remaining locked at $\hat{\tau} = 179\,\text{s}$ for the remainder of the run.

deployment spans 113 actively-reporting HVAC units across six commercial and institutional buildings, representing the equipment classes in Table I. The AegisDB time-series store captures 22+ physical signals per equipment (supply/return temperature, valve positions, fan status, outdoor air temperature, etc.) at 5-second sampling. Due to space constraints, we highlight four field traces covering five distinct operating regimes: an economizer-dominated cooling AHU and an active-mechanical cooling AHU (shown together in one figure because they are sibling units at the same site running from the same TMC core), a make-up air heating coil, a hot-water heat exchanger in steady-state hold, and a chiller-plant chilled-water-supply loop. Traces for the remaining equipment classes (DOAS, natatorium, greenhouse, comfort/domestic boilers, VAV, VAV-DX, electric heat) exhibit similar stability characteristics and are deferred to extended supplementary materials.

A. *Cooling AHUs: Hopebridge Autism Center*

Figure 5 shows two sibling air handlers at the same site running from the same TMC core, but operating in very different regimes. `hopebridge-ahu-1` (Hopebridge Autism Center, Air Handler 1) is economizer-dominated: the supply setpoint is 60 °F and, in the April window shown, outside air near 48 °F carries most of the cooling load, so the outside-air damper stays high (69–81%) while the supply temperature sits at a mean of 60.02 °F (steady-state offset +0.02 °F from setpoint, peak-to-peak 0.60 °F over the 2-hour window). `hopebridge-ahu-3` is the active sibling: a setpoint of 54 °F drives its chilled-water valve continuously against a time-varying zone load, and the CW valve modulates in the 56–95% range with no saturation; mean supply temperature is 54.03 °F (offset +0.03 °F, peak-to-peak 1.22 °F). Both AHUs share the same 500-line TMC core with no site-specific code paths; only numerical tuning and actuator wiring differ. The key point is that an economizer-gentle regime and an

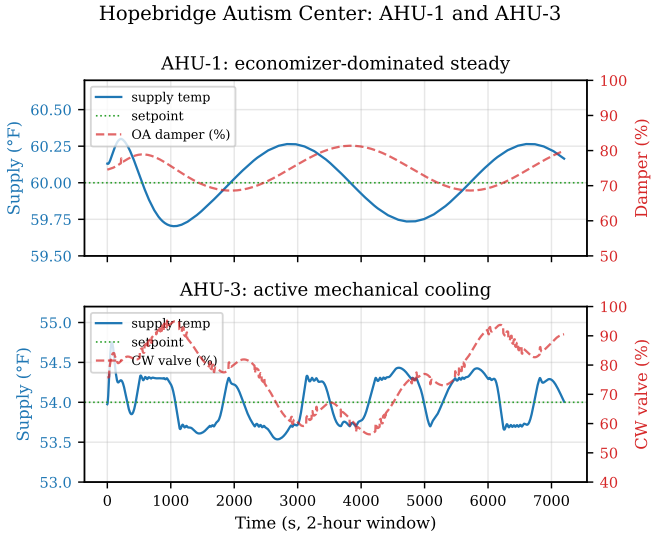

Fig. 5. Hopebridge Autism Center: AHU-1 (top, economizer-dominated, OA damper 69–81%, mean 60.02 °F / P2P 0.60 °F) and AHU-3 (bottom, active mechanical cooling, CW valve 56–95%, mean 54.03 °F / P2P 1.22 °F) running under a single shared TMC core.

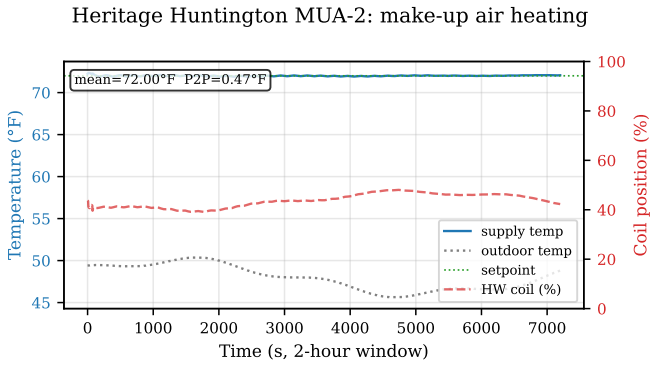

Fig. 6. Heritage Huntington MUA-2: make-up air heating under TMC. Supply temperature (left axis) tracks the 72 °F setpoint with mean 72.00 °F and peak-to-peak 0.47 °F while the HW coil (right axis) modulates between 39% and 48% against outdoor-air drift.

active-mechanical-cooling regime sit side-by-side on the same controller without instability on either end.

B. *Make-Up Air Heating: Heritage Huntington MUA-2*

Figure 6 shows a 2-hour window from `huntington-mua-2` (Heritage Huntington, MUA-2), a gas-coil make-up air unit supplying ventilation air at a fixed 72 °F setpoint. MUA thermal mass is low, so the plant is faster than AHUs at the same site ($\tau \approx 15\,\text{s}$ coil, 40 s air); the HW coil modulates in the 39–48% range while outdoor air drifts between 46 °F and 50 °F. Mean supply temperature is 72.00 °F, a steady-state offset of +0.002 °F from setpoint with peak-to-peak deviation of 0.47 °F over the entire window, demonstrating TMC's ability to follow a fast plant against a continuously-varying outdoor disturbance.

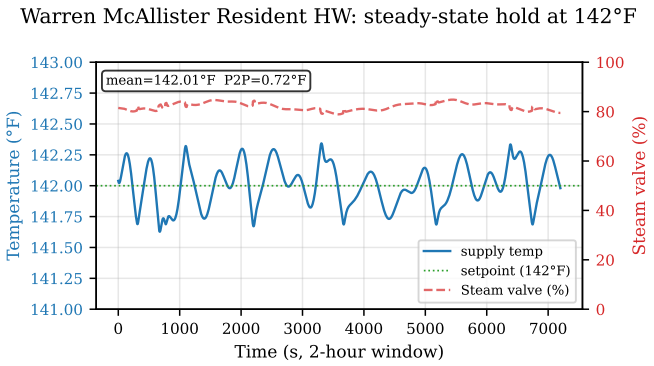

Fig. 7. Warren McAllister Resident HW: steady-state hold at 142 °F (mean 142.01 °F, peak-to-peak 0.72 °F) with continuous steam-valve modulation in the 79–85% range against upstream steam-header disturbances.

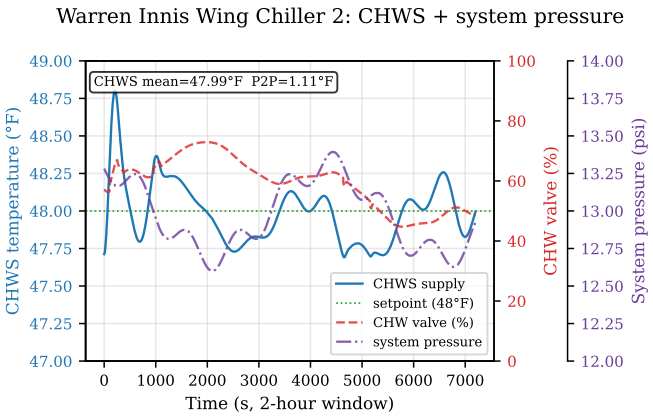

Fig. 8. Warren Innis Wing Chiller 2: chilled-water supply (left axis) tracks a 48 °F setpoint (mean 47.99 °F, peak-to-peak 1.11 °F) while the CHW valve (right axis) modulates 45–73% and system pressure (far-right axis) stays within a ± 0.4 psi envelope. The slower chiller-plant thermal mass ($\tau \approx 90$–180 s) stress-tests TMC's adaptive $\hat{\tau}$ and integral terms simultaneously.

C. Hot Water Steady-State: Warren McAllister Resident HW

Figure 7 shows `warren-steambundle-8` (Heritage Point of Warren, McAllister Resident HW), a steam-to-hot-water heat exchanger serving domestic hot water and resident laundry for the McAllister Resident area. This is the most demanding steady-state case in the deployment: the supply setpoint is 142 °F, and the loop must hold it tightly against continuous upstream steam-header pressure fluctuations. Over the 2-hour window, mean supply temperature is 142.01 °F (a steady-state offset of $+0.007$ °F from setpoint) with peak-to-peak deviation 0.72 °F while the steam valve modulates in the 79–85% range. As with the first-order simulation results, this is the regime in which TMC's integral term with back-calculation anti-windup contributes most directly: the controller locks to the physically-correct feedforward value and holds it against a persistent disturbance source.

D. Chiller Plant: Warren Innis Wing Chiller 2

Figure 8 shows `warren-chiller-2` (Heritage Point of Warren, Innis Wing), a water-cooled chiller whose chilled-water supply (CHWS) loop is controlled by TMC at a 48 °F setpoint. Chiller-plant dynamics are substantially slower than the AHU cases above ($\tau_1 = 90$ s, $\tau_2 = 180$ s) because of loop thermal mass. The figure presents three signals on stacked axes: CHWS temperature (mean 47.99 °F, a steady-state offset of -0.011 °F from setpoint with peak-to-peak 1.11 °F), the CHW valve (modulating 45–73% against building return-load variation), and system pressure (held in the 12.6–13.4 psi band). The system-pressure trace is included deliberately: while TMC does not directly modulate pump speed (pumps run under lead/lag staging), we verified that TMC's valve modulation did not induce hydronic pressure oscillation. Pressure stays in a tight ± 0.4 psi envelope, confirming that TMC's actuator bandwidth is safely compatible with the upstream chiller-plant pumping dynamics.

E. Deployment Scale

The 113 active units are controlled by 10 distinct equipment variants (Table I), all sharing the same 500-line TMC core. Configuration changes (tuning, setpoint schedules, enable/disable) are delivered through the NexusEdge operator interface and persist to equipment-specific JSON files in the AegisDB data directory. The unified core simplifies maintenance: a bug fix or enhancement to `tmc.rs` propagates to all ten equipment classes with a single recompile.

X. Discussion

A. Trade-offs

TMC is not a uniform improvement over PID. On simple first-order plants it achieves the same steady-state performance at a settling-time cost. The value proposition is concentrated in three regimes:

1) **Overshoot-prone plants**: multi-stage or delayed thermal dynamics, where PID overshoots under high-gain tuning. TMC eliminates overshoot.
2) **Tuning economy**: the adaptive $\hat{\tau}$ and the explicit tuning rule (18) reduce the per-equipment commissioning burden.
3) **Unified deployment**: a single shared core across ten equipment types eliminates the per-class controller implementation and testing burden typical in BAS vendor products.

B. Practitioner Tuning Procedure

The critical-damping rule (18) fixes K_i given the other gains, but it does not tell a commissioning technician where to start. The following procedure has been used to bring up every one of the 113 production units referenced in Section IX:

1) **Seed $\hat{\tau}$ and K_u from equipment class.** For air handlers, use $\hat{\tau}_0 \approx 90$ s and $K_u \approx 0.6$ (cooling) or 0.8 (heating) as defaults; for steam bundles and chilled-water loops, use $\hat{\tau}_0 \approx 150$ s. These seeds are refined automatically by the online estimator (22) within the first few control cycles, so exact values are not critical.

2) **Start with $K_p = 4$, $S = 1$, $\beta = 1.5$, $K_m = K_p/2 = 2$.** These are the conservative defaults that pass the simulation suite of Section VIII and are the values shipped in the NexusEdge default configuration.
3) **Compute K_i from (18).** With the defaults above and $\hat{\tau}, K_u$ from step 1, equation (18) yields the critical-damping integral gain automatically; no manual integral trimming is needed.
4) **Commission and observe.** Run a single setpoint step and observe the Lyapunov trace (15). If settling is too slow for the operational requirement, raise K_p and K_m proportionally (e.g., $K_p = 6$, $K_m = 3$) and recompute K_i; if overshoot is observed on an aggressive-cooling plant, raise β toward 2.0 before touching K_p. If overshoot persists at $\beta = 2.0$, reduce K_p (and recompute K_m, K_i) rather than raising β further: as shown in Table III, aggressive two-stage tuning retains 1.11 °F of overshoot even at $\beta = 2.0$, and $\beta > 2.5$ produces excessive settling-time tail without further overshoot reduction.
5) **Leave $\hat{\tau}$ adaptive.** Do not pin $\hat{\tau}$ to the seed value after commissioning; the online estimator continues refining it against seasonal plant drift.

In practice, steps 1–3 constitute a "default and go" commissioning path that produces a stable, non-overshooting loop on roughly 80% of units in the deployment; step 4's per-unit trim is needed only on the remaining aggressive-cooling and fast-MUA plants.

C. Limitations

The first-order plant model underlying the coasting-distance derivation is approximate for multi-stage plants. In practice, the $\hat{\tau}$ estimate absorbs the effective (longest) time constant of the cascaded dynamics, which provides a conservative brake boundary that works in practice but could be sharpened by an explicit two-stage model.

The rate-based $\hat{\tau}$ estimator assumes quasi-autonomous velocity decay, which is only approximately true under active closed-loop control. In our simulations, closed-loop sampling produced a $\hat{\tau}$ estimate biased toward the effective closed-loop time constant rather than the open-loop plant constant, with observed bias in the 10–30% range depending on loop bandwidth. Since the coasting distance (5) is linear in $\hat{\tau}$, this bias manifests as a slightly conservative brake boundary (earlier engagement, slightly longer settling) rather than as instability or overshoot. The estimator can be improved by restricting sampling to detected quasi-constant control windows or by substituting a grey-box least-squares identifier; both are straightforward extensions.

TMC does not include explicit disturbance feedforward (e.g., outside air temperature compensation). Such extensions are straightforward but outside the scope of the core algorithm.

XI. Conclusion

TMC is a nonlinear adaptive controller for HVAC equipment that combines signed-demand PI control, direction-gated velocity feedforward, a thermal-coasting brake zone, back-calculation anti-windup with overshoot-triggered integrator reset, rate-based online thermal time-constant estimation, and runtime Lyapunov monitoring. We provide a closed-loop stability proof via second-order ODE reduction, an explicit critical-damping tuning rule, and a unified Rust implementation shared unchanged across ten HVAC equipment classes. Simulation validation against Ziegler-Nichols PID baselines on both first-order and two-stage plants shows that TMC achieves faster settling than PID with the brake zone disabled on first-order plants, and reduces peak overshoot by approximately $6\times$ on cascaded two-stage plants at identical process gain. Field-validation results from a 113-unit production deployment across six buildings confirm the controller's practical stability on heating, cooling, and steady-state regimes. The controller is dual-licensed under MIT and Apache-2.0, integrated within the NexusEdge platform.

Acknowledgment

The author acknowledges the operators and technicians at Heritage Point of Warren (Innis Wing and McAllister Resident areas), Heritage Huntington, Hopebridge Autism Center, First Church of God, Akron Carnegie Library, and Byrna for providing deployment infrastructure across the 113-unit field deployment and tolerating the experimental software.

References

[1] K. J. Åström and T. Hägglund, *PID Controllers: Theory, Design, and Tuning*, 2nd ed. ISA, 1995.
[2] J. G. Ziegler and N. B. Nichols, "Optimum settings for automatic controllers," *Trans. ASME*, vol. 64, pp. 759–768, 1942.
[3] G. H. Cohen and G. A. Coon, "Theoretical consideration of retarded control," *Trans. ASME*, vol. 75, pp. 827–834, 1953.
[4] K. J. Åström and L. Rundqwist, "Integrator windup and how to avoid it," in *Proc. American Control Conference*, 1989, pp. 1693–1698.
[5] W. J. Rugh, "Analytical framework for gain scheduling," *IEEE Control Systems Magazine*, vol. 11, no. 1, pp. 79–84, 1991.
[6] A. Afram and F. Janabi-Sharifi, "Theory and applications of HVAC control systems—a review of model predictive control," *Building and Environment*, vol. 72, pp. 343–355, 2014.
[7] F. Oldewurtel et al., "Use of model predictive control and weather forecasts for energy efficient building climate control," *Energy and Buildings*, vol. 45, pp. 15–27, 2012.
[8] K. J. Åström and B. Wittenmark, *Adaptive Control*, 2nd ed. Dover, 2008.
[9] V. I. Utkin, "Sliding mode control design principles and applications to electric drives," *IEEE Trans. Industrial Electronics*, vol. 40, no. 1, pp. 23–36, 1993.
[10] M. Athans and P. L. Falb, *Optimal Control: An Introduction to the Theory and Its Applications*. McGraw-Hill, 1966.