INTERNATIONAL JOURNAL OF LATEST TECHNOLOGY IN ENGINEERING, MANAGEMENT & APPLIED SCIENCE (IJLTEMAS)

ISSN 2278-2540 | DOI: 10.51583/IJLTEMAS | Volume XV, Issue VII, July 2026

Core Temperature Estimation and Thermal Management of Lithium-Ion Batteries Using a Nonlinear Adaptive Extended Kalman Filter

Kalpesh Madhav Mahajan[1], Suhas Manohar Shembekar[2], Vijay Mangalsing Deshmukh[3]

[1, 2, 3] Department of Electrical Engineering, SSBT's College of Engineering and Technology, Bambhori, Jalgaon, Maharashtra 425001, India

DOI: https://doi.org/10.51583/IJLTEMAS.2026.150700091

Received: 29 July 2026; Accepted: 03 August 2026; Published: 14 August 2026

ABSTRACT

The core temperature of a cylindrical lithium-ion cell cannot be measured directly, yet it governs both safety and ageing. This paper presents a nonlinear adaptive extended Kalman filter (AEKF) that estimates core temperature from a single surface thermistor, together with a battery thermal management strategy driven by the resulting estimate. Three features distinguish the formulation from earlier two-node observers. First, heat generation is evaluated at the estimated core temperature rather than treated as an exogenous input; because internal resistance follows an Arrhenius law, this makes the process model genuinely nonlinear in the state and requires an explicit Jacobian, derived here in closed form. The state-feedback term reaches 25.4 % of the dominant conduction term at peak current, and substituting the measured surface temperature into the heat-generation expression incurs an error of up to 1.35 W. Second, the lumped two-node parameters are identified against a radially resolved finite-volume reference model using an excitation containing coolant-flow steps; omitting those steps yields a surface capacitance an order of magnitude too small and an observer that mispredicts the core whenever the pump switches. Third, innovation-based covariance adaptation is made safe for closed-loop use by a Student-t test on the innovation mean that distinguishes sensor degradation from transient model error, and by bounded directional process-noise inflation applied only to the measured state. Validation is by simulation only. Against a structurally different reference plant with deliberate parameter mismatch and a mid-run doubling of thermistor noise, the proposed filter attains a core-temperature RMSE of 0.348 degrees Celsius, against 0.359 for a fixed-gain EKF, 0.429 for a Sage-Husa AEKF and 0.447 for the linear formulation in which heat generation is an input. Under twenty Monte-Carlo realisations with random parameter error the proposed filter and the fixed-gain EKF are statistically indistinguishable. In closed loop, estimate-driven cooling holds the same peak core temperature as a conventional surface-triggered controller while consuming 47.5 % less pump energy.

Keywords: lithium-ion battery; core temperature estimation; extended Kalman filter; adaptive filtering; battery thermal management

INTRODUCTION

Lithium-ion cells degrade faster and fail more abruptly at elevated temperature. The quantity that matters is the temperature inside the jellyroll, not the temperature of the casing, and in a cylindrical 21700 cell under a fast-charge load the two can differ by several degrees. Because instrumenting the core of a production cell is impractical, the core temperature must be inferred from surface measurements through a thermal model and a state observer.

Two-node lumped thermal models coupled to Kalman-type observers are the standard approach [1]–[3]. This paper revisits that approach and addresses three weaknesses that recur in the literature and that were identified in review of an earlier version of this work.

The justification for an extended Kalman filter is usually absent. If heat generation is treated as a measured input, the two-node model is linear and time-invariant, no Jacobian exists, and the word

ISSN 2278-2540 | DOI: 10.51583/IJLTEMAS | Volume XV, Issue VII, July 2026

"extended" is not earned. This objection is correct as far as it goes. The resolution adopted here is not to rename the filter but to remove the modelling shortcut that produced the linearity: heat generation depends on the internal resistance, the internal resistance depends on temperature through an Arrhenius law, and the relevant temperature is the core temperature, which is precisely the unmeasurable state. Treating heat generation as an input requires substituting the surface temperature into that expression, which is an approximation with a quantifiable cost. Section III derives the nonlinear model and its Jacobian and quantifies both the strength of the nonlinearity and the error incurred by ignoring it.

The novelty of "adaptive" covariance schemes is frequently overstated. Innovation-based covariance matching dates to Mehra [4] and Mohamed and Schwarz [5]; the Sage–Husa recursion [6] is equally well established. Combining one of these with a two-node thermal model is not by itself a contribution. Section IV therefore states precisely what is and is not claimed here, and Section VI reports a result that does not favour naive adaptation: a Sage–Husa AEKF is worse than a well-tuned fixed-gain EKF on this problem, and considerably less reliable across parameter realisations.

Reported accuracies are rarely referred to a model floor. An observer built on a lumped model cannot outperform the model's own structural error. Section VI reports that floor explicitly and interprets the estimator results against it.

The contributions are therefore: a nonlinear two-node formulation with a closed-form Jacobian and a quantitative justification for it (Section III); an identification procedure that exposes and corrects a failure mode of conventional lumped-parameter fitting (Section V-B); a safeguarded adaptation mechanism whose components are individually motivated and whose benefit is honestly bounded (Section IV); and a closed-loop thermal management comparison performed at matched thermal performance so that the energy comparison is meaningful (Section VII).

Scope and limitations

All results in this paper are obtained by numerical simulation. No cell was instrumented, no hardware-in-the-loop test was performed, and no microcontroller execution time was measured. The reference plant against which the observers are validated is a radially resolved finite-volume model, structurally different from the lumped model used by the observers, so the reported errors include genuine structural model error rather than only sensor noise. This is stronger evidence than a self-consistency check, but it is not experimental validation, and every quantitative claim below should be read as a model prediction. Section VIII lists the specific claims that require experimental confirmation before they can be relied upon.

Reference Plant and Operating Scenario

Cell and thermal domains

The reference plant resolves three thermal domains for a 21700 cylindrical NMC cell.

The **jellyroll** is treated as a solid cylinder with uniform volumetric heat generation and radial conduction,

$$\rho c_p \frac{\partial T}{\partial t} = \frac{1}{r}\frac{\partial}{\partial r}\left(r\, k_r \frac{\partial T}{\partial r}\right) + \dot{q}'''(t, T)$$

discretised conservatively on $N_r = 20$ annular finite volumes and integrated with an explicit sub-step of 0.05 s, which satisfies the stability limit $\Delta t < \Delta r^2/(2\alpha)$ with a margin of approximately four.

The casing is a lumped steel node coupled to the outer jellyroll volume through a contact conductance G_c. Its temperature is what a thermistor bonded to the can physically reads, and it is the only measured state.

The cold plate is a lumped node with finite thermal mass, comprising the per-cell share of the aluminium plate and its coolant hold-up:

ISSN 2278-2540 | DOI: 10.51583/IJLTEMAS | Volume XV, Issue VII, July 2026

$$C_p \frac{dT_p}{dt} = G_{sp}(T_s - T_p) - G_{pa}(T_p - T_a) - UA_{eff}(\dot{m})(T_p - T_f)$$

Modelling the plate as a finite mass rather than an infinite sink is essential. With an infinite sink the pump-off condition still removes heat, which understates the thermal problem substantially; with a finite plate, switching the pump off leaves the plate to heat up with the cell, which is what physically happens.

The plate-to-coolant conductance follows an ε–NTU counter-flow model with a Dittus–Boelter flow dependence,

$$UA_{eff}(\dot{m}) = \varepsilon\, \dot{m} c_f, \qquad \varepsilon = 1 - \exp\left(-\frac{UA(\dot{m})}{\dot{m} c_f}\right), \qquad UA(\dot{m}) = UA_{ref}\left(\frac{\dot{m}}{\dot{m}_{max}}\right)^{0.8}$$

with $UA_{eff} = 0$ when the pump is off.

Electro-thermal coupling

Heat generation comprises an irreversible and a reversible term,

$$\dot{Q}_{gen}(I, T) = I^2 R_{int}(T) + I\,(T + 273.15)\,\frac{\partial U_{ocv}}{\partial T}$$

with an Arrhenius internal resistance

$$R_{int}(T) = R_{ref} \exp\left[\frac{E_a}{R_g}\left(\frac{1}{T + 273.15} - \frac{1}{T_{ref}}\right)\right]$$

Over the range 25 °C to 45 °C the internal resistance falls from 22.0 mΩ to 14.4 mΩ, a variation of 34.4 %. This dependence is the origin of the state nonlinearity analysed in Section III.

Duty profile

The excitation is a composite duty representing hot-climate operation: urban driving, followed by DC fast charging, followed by rest.

The speed trace is a 1369 s urban schedule synthesised from a fixed micro-trip template and rescaled to reproduce the published statistics of the EPA Urban Dynamometer Driving Schedule. The generated trace matches the reference values closely, as shown in Table I. It is not the exact EPA time series, and this is stated as a limitation; substituting the certified trace would be a straightforward change.

Table I — Generated drive cycle versus published UDDS statistics

Quantity	Published UDDS	Generated trace
Duration (s)	1369	1369
Distance (km)	12.07	12.07
Maximum speed (km/h)	91.2	91.2
Mean speed (km/h)	31.5	31.7
Number of stops	17	17
Idle fraction	≈ 0.19	0.21

ISSN 2278-2540 | DOI: 10.51583/IJLTEMAS | Volume XV, Issue VII, July 2026

The conversion from speed to cell current is stated explicitly, since it was not previously documented. Longitudinal dynamics give the tractive force

$$F_{trac} = m_v a + 1/2\,\rho_{air} C_d A_f v^2 + m_v g C_{rr}$$

and the battery power follows with separate traction and regeneration efficiencies,

$$P_{batt} = \begin{cases} F_{trac} v/\eta_{dt} + P_{aux}, & F_{trac} v \geq 0 \\ F_{trac} v\,\eta_{regen} + P_{aux}, & F_{trac} v < 0 \end{cases}$$

$$I_{cell} = \frac{P_{batt}}{N_s V_{nom} N_p}$$

Vehicle and pack parameters are listed in Table II. The auxiliary load of 1500 W represents cabin air-conditioning under a 40 °C ambient.

The driving phase is followed by a 3 C constant-current DC fast charge of 1000 s and a 600 s rest. The fast-charge phase is included because a single cell in a correctly sized pack dissipates well under 0.2 W on the UDDS alone, which is not by itself a thermal-management problem. It is important to state this plainly: on urban driving at moderate ambient, a properly sized pack does not require active cooling, and a paper that claims otherwise is not describing a real vehicle. The thermal problem in hot climates arises from the combination of high ambient temperature and fast charging.

Consistent with production practice, the coolant loop rejects heat through the radiator at 30 °C during driving and the refrigerant chiller is engaged during fast charging, stepping the coolant inlet to 22 °C over a 20 s ramp. This step is also the most demanding condition for a lumped observer, because it imposes a fast boundary transient on top of the current transient.

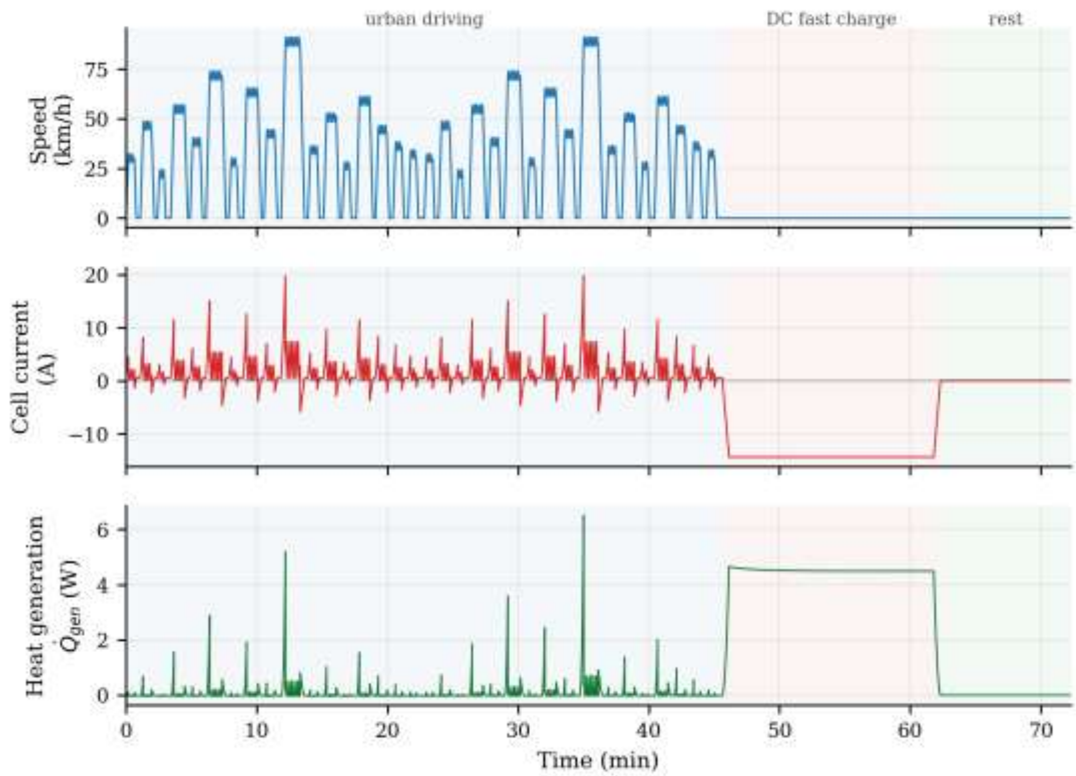

Fig. 2. Composite duty profile: vehicle speed, cell current and resulting heat generation. Shading marks the urban-driving, DC-fast-charge and rest phases.

ISSN 2278-2540 | DOI: 10.51583/IJLTEMAS | Volume XV, Issue VII, July 2026

Table II — Principal parameters

Symbol	Description	Value
R_{cell}, H_{cell}	cell radius, height	10.50 mm, 70.0 mm
m_{cell}, c_p	cell mass, specific heat	70.0 g, 1000 J kg^{-1} K^{-1}
k_r	radial conductivity	0.90 W m^{-1} K^{-1}
Q_{nom}, V_{nom}	cell capacity, voltage	4.80 Ah, 3.60 V
R_{ref}, E_a/R_g	resistance at 25 °C, activation temp.	22.0 mΩ, 2000 K
$\partial U/\partial T$	entropic coefficient	-1.50×10^{-4} V K^{-1}
G_c	jellyroll–casing conductance	12.0 W K^{-1}
G_{sp}	casing–plate conductance	0.717 W K^{-1}
C_p, G_{pa}	plate capacity, plate–ambient	70.0 J K^{-1}, 0.150 W K^{-1}
UA_{ref}	plate–coolant UA at $\dot{m}_{max}$	6.0 W K^{-1}
m_v, C_dA_f, C_{rr}	vehicle mass, drag, rolling	1250 kg, 0.29 × 2.10 m², 0.010
$N_s \times N_p$	pack configuration	96 × 8
P_{aux}	auxiliary load	1500 W
T_a, T_f	ambient, coolant inlet (radiator / chilled)	40 °C, 30 / 22 °C
T_{safe}	safe core temperature limit	45 °C
$\dot{m}_{max}$, τ_{pump}	max flow, pump time constant	8.0 L min^{-1}, 2.0 s
σ_T	thermistor noise (phase 1 / 2)	0.30 / 0.60 °C
Δt	estimator sampling period	0.1 s

Nonlinear Two-Node Model and Its Jacobian

State-space formulation

The observer uses a lumped two-node model with state $\mathbf{x} = [T_c, T_s]^\top$, input $\mathbf{u} = [I, T_a, T_p]^\top$ and output $z = T_s$:

$$\frac{dT_c}{dt} = \frac{1}{C_c}\left[\dot{Q}_{gen}(I, T_c) - \frac{T_c - T_s}{R_{cs}}\right]$$

$$\frac{dT_s}{dt} = \frac{1}{C_s}\left[\frac{T_c - T_s}{R_{cs}} - \frac{T_s - T_a}{R_a} - G_{sp}(T_s - T_p)\right]$$

ISSN 2278-2540 | DOI: 10.51583/IJLTEMAS | Volume XV, Issue VII, July 2026

$$z = \mathbf{H}\mathbf{x} + v, \qquad \mathbf{H} = [0 \quad 1]$$

The plate temperature T_p enters as a measured input. Instrumenting a cold plate is routine in liquid-cooled packs, and taking T_p as a measurement rather than a state keeps the observer second-order and removes any need to model pump hydraulics inside the filter.

Why the model is nonlinear

Substituting the heat-generation expression into the first state equation gives

$$\frac{dT_c}{dt} = \frac{1}{C_c}\left[I^2 R_{ref} \exp\left(\frac{E_a}{R_g}\left(\frac{1}{T_c + 273.15} - \frac{1}{T_{ref}}\right)\right) + I(T_c + 273.15)\frac{\partial U}{\partial T} - \frac{T_c - T_s}{R_{cs}}\right]$$

which is manifestly nonlinear in T_c. The nonlinearity is not decorative. Differentiating,

$$\frac{\partial \dot{Q}_{gen}}{\partial T_c} = I^2 \frac{dR_{int}}{dT_c} + I\frac{\partial U}{\partial T}, \qquad \frac{dR_{int}}{dT_c} = -R_{int}(T_c)\frac{E_a/R_g}{(T_c + 273.15)^2}$$

The natural dimensionless measure of the nonlinearity is the ratio of this state-feedback term to the conduction term $1/R_{cs}$ that dominates the same equation, that is $\left|\partial \dot{Q}_{gen}/\partial T_c\right| R_{cs}$. Over the operating envelope this ratio reaches 25.4 % at peak current, and it is therefore not negligible at the accuracy sought.

An alternative and more direct argument concerns the cost of the linear formulation. If $\dot{Q}_{gen}$ is treated as an exogenous input it must be evaluated at some known temperature, and the only available one is the measured surface temperature. The resulting heat-generation error is

$$\delta\dot{Q} \approx \frac{\partial \dot{Q}_{gen}}{\partial T}(T_c - T_s)$$

which grows with exactly the gradient the observer exists to resolve. Over the operating envelope this reaches 1.35 W, against a mean generation of order 1 W. The linear formulation is therefore biased in precisely the regime that matters, and Section VI-C confirms this empirically: its RMSE is 24 % higher than that of the nonlinear filter, and its bias is nearly three times larger.

Jacobian and discretisation

The continuous-time Jacobian is available in closed form:

$$\mathbf{J}_c(\mathbf{x}, \mathbf{u}) = \begin{bmatrix} \frac{1}{C_c}\left(\frac{\partial \dot{Q}_{gen}}{\partial T_c} - \frac{1}{R_{cs}}\right) & \frac{1}{C_c R_{cs}} \\ \frac{1}{C_s R_{cs}} & \frac{1}{C_s}\left(-\frac{1}{R_{cs}} - \frac{1}{R_a} - G_{sp}\right) \end{bmatrix}$$

The prediction step integrates the nonlinear dynamics with a fourth-order Runge–Kutta scheme using two sub-steps per sample, and the discrete Jacobian is the exact exponential of the frozen linearisation,

$$\mathbf{F}_k = \exp(\mathbf{J}_c(\hat{\mathbf{x}}_{k-1}, \mathbf{u}_{k-1})\,\Delta t)$$

evaluated in closed form for the 2×2 case,

$$\exp(\mathbf{A}) = e^{\tau}\left[\cosh(\delta)\mathbf{I} + \frac{\sinh(\delta)}{\delta}(\mathbf{A} - \tau\mathbf{I})\right], \quad \tau = 1/2\,\mathrm{tr}\mathbf{A}, \quad \delta = \sqrt{\tau^2 - \det\mathbf{A}}$$

The measurement map is linear, so no measurement Jacobian is required; all nonlinearity resides in the process model. This is stated explicitly because the absence of a measurement Jacobian has previously been read as evidence that no Jacobian exists at all.

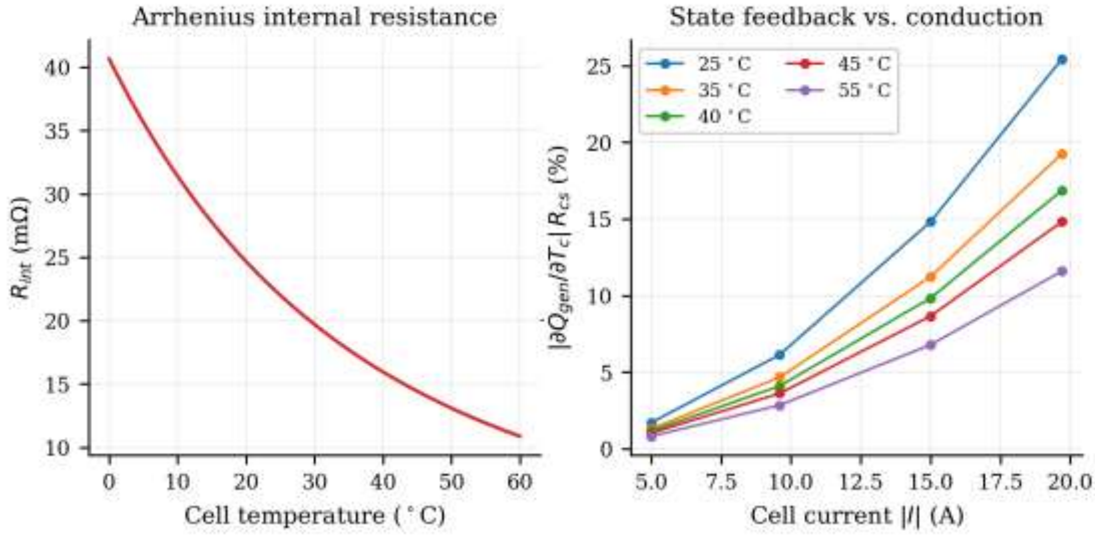

Fig. 1. Evidence that the plant is nonlinear in the state. Left: Arrhenius internal resistance, which falls 34.4 % between 25 °C and 45 °C. Right: the state-feedback term $\left|\partial \dot{Q}_{gen}/\partial T_c\right| R_{cs}$ expressed as a percentage of the conduction term, reaching 25.4 % at peak current.

Observability

The observability matrix of the pair $(\mathbf{F}_\mathrm{k}, \mathbf{H})$ is

$$\mathcal{O}_\mathrm{k} = \begin{bmatrix} \mathbf{H} \\ \mathbf{H}\mathbf{F}_\mathrm{k} \end{bmatrix} = \begin{bmatrix} 0 & 1 \\ \mathrm{F}_{21} & \mathrm{F}_{22} \end{bmatrix}$$

which has rank 2 whenever $\mathrm{F}_{21} \neq 0$, that is whenever R_{cs} is finite. At a representative operating point the rank is 2 and the condition number is 594. The core temperature is therefore observable from the surface measurement, but the pair is poorly conditioned, which explains why the achievable accuracy is limited by model quality rather than by sensor noise — a point that Section VI-B makes quantitative.

Proposed Adaptive Filter

What is and is not claimed

The measurement- and process-noise covariances of a Kalman filter are rarely known and rarely constant. Estimating them from the innovation sequence is a classical idea: Mehra [4] introduced innovation correlation methods, Mohamed and Schwarz [5] developed the covariance-matching form used here, and Sage and Husa [6] gave the exponentially weighted recursion. None of this is claimed as novel**.** What is proposed is a set of safeguards that make covariance matching usable inside a closed thermal control loop, together with the evidence in Section VI that without them adaptation actively degrades performance on this problem.

The safeguards address a specific failure. Two quite different faults inflate the innovation sequence:

Sensor degradation The thermistor becomes noisier. Innovations grow but remain zero-mean. The correct response is a permanent increase in the measurement-noise estimate.

Transient model error When the pump switches on or the chiller engages, the lumped model — which has no radial transport delay - mispredicts how fast the casing cools. Innovations grow and acquire a bias. Here

ISSN 2278-2540 | DOI: 10.51583/IJLTEMAS | Volume XV, Issue VII, July 2026

the surface measurement briefly stops being informative about the core: the model drags the core estimate along with the measured surface, while the true core, buffered by the outer jellyroll, has barely moved. The correct response is a temporary, reversible loss of confidence.

Scalar covariance matching cannot separate these, and raising R — the Sage–Husa response — de-weights the measurement globally, degrading the surface estimate as well.

Bias discrimination

The discriminator is a Student-t statistic on the mean of the innovation window $\{\varepsilon_i\}_{i=k-L+1}^{k}$:

$$\tau_k = \frac{|\bar{\varepsilon}_k|}{s_{\varepsilon,k}/\sqrt{L}}, \qquad \bar{\varepsilon}_k = \frac{1}{L}\sum_{i=k-L+1}^{k} \varepsilon_i\,, \qquad s_{\varepsilon,k}^2 = \frac{1}{L}\sum_{i=k-L+1}^{k} (\varepsilon_i - \bar{\varepsilon}_k)^2$$

A zero-mean innovation sequence gives τ_k of order unity irrespective of its variance, whereas a biased sequence gives a large value. In the runs reported below τ_k remains below 1 during healthy operation and exceeds 5 at the chiller transient, so the two regimes separate cleanly.

An autocorrelation-based whiteness test was also evaluated and rejected: at a 0.1 s sampling period the innovations are correlated even when the filter is healthy, because the surface time constant spans many samples, so the whiteness test flags model error almost continuously.

Slow loop — sensor statistics

When $\tau_k \le \tau_{crit}$ the innovations are consistent with a zero-mean sequence and the noise statistics are updated with exponential forgetting:

$$\hat{R}_{s,k} = \Pi_{[R_{min},R_{max}]}\left\{\lambda\hat{R}_{s,k-1} + (1-\lambda)\left[s_{\varepsilon,k}^2 - \frac{1}{L}\sum_i \mathbf{H}\,\mathbf{P}_i^-\mathbf{H}^\top\right]\right\}$$

$$\hat{\mathbf{Q}}_k = \Pi_{[Q_{min},Q_{max}]}\left\{\lambda\hat{\mathbf{Q}}_{k-1} + \frac{1-\lambda}{L}\sum_i (\mathbf{K}_i\varepsilon_i)\,(\mathbf{K}_i\varepsilon_i)^\top\right\}$$

where Π denotes projection: a scalar clip for R and an eigenvalue clip for **Q**,

$$\Pi_{[a,b]}\{\mathbf{Q}\} = \mathbf{V}\mathrm{diag}\left(\min\left(\max\left(w_j, a\right), b\right)\right)\mathbf{V}^\top, \qquad \mathbf{Q} = \mathbf{V}\mathrm{diag}\left(w_j\right)\mathbf{V}^\top$$

Fast loop — directional inflation

When $\tau_k > \tau_{crit}$ the innovations are biased, the noise estimates are frozen, and a bounded inflation is applied to the process noise of the **measured** state only:

$$\mathbf{Q}_{eff,k} = \hat{\mathbf{Q}}_k + \mathrm{diag}(0,\ q_k), \qquad q_k = \min\left(\kappa\,\bar{\varepsilon}_k^2,\ q_{max}\right)$$

The direction matters. With $\mathbf{H} = [0\ \ 1]$ the gain is

$$\mathbf{K}_k = \frac{1}{P_{22}^- + q_k + R}\begin{bmatrix} P_{12}^- \\ P_{22}^- + q_k \end{bmatrix}$$

so inflating q_k reduces the core gain K_1 while increasing the surface gain K_2. The core estimate is decoupled from the momentarily misleading measurement while the surface estimate remains accurate. A scalar inflation of R reduces both gains and cannot achieve this; the surface RMSE penalty this imposes on the Sage–Husa filter is visible in Table V.

ISSN 2278-2540 | DOI: 10.51583/IJLTEMAS | Volume XV, Issue VII, July 2026

Lemma 1 (bounded gain). With $\mathrm{R} \geq \mathrm{R}_{\min} > 0$, $\mathrm{q}_{\mathrm{k}} \geq 0$ and $\mathbf{Q}$ projected onto $[\mathrm{Q}_{\min}, \mathrm{Q}_{\max}]\mathbf{I}$, the gain satisfies

$$\|\mathbf{K}_{\mathrm{k}}\| \leq \frac{\|\mathbf{P}_{\mathrm{k}}^{-}\mathbf{H}^{\top}\|}{\mathrm{R}_{\min}}$$

and $\mathbf{P}_{\mathrm{k}}^{-}$ is bounded above because $\mathbf{F}_{\mathrm{k}}$ is a matrix exponential of a Hurwitz matrix with bounded entries and $\mathbf{Q}_{\mathrm{eff,k}} \preccurlyeq (\mathrm{Q}_{\max} + \mathrm{q}_{\max})\mathbf{I}$. The adapted filter therefore cannot diverge, and with $\kappa = 0$ and adaptation disabled it reduces exactly to the EKF.

Algorithm

Algorithm 1 — Bias-discriminated adaptive EKF

Inputs: z_{k}, $\mathbf{u}_{\mathrm{k}} = [\mathrm{I}_{\mathrm{k}}, \mathrm{T}_{\mathrm{a,k}}, \mathrm{T}_{\mathrm{p,k}}]^{\top}$. State: $\hat{\mathbf{x}}_{\mathrm{k}-1}$, $\mathbf{P}_{\mathrm{k}-1}$, $\widehat{\mathbf{Q}}_{\mathrm{k}-1}$, $\widehat{\mathrm{R}}_{\mathrm{s,k}-1}$, $\mathrm{q}_{\mathrm{k}-1}$, buffers $\{\varepsilon\}$, $\{\mathbf{H}\mathbf{P}^{-}\mathbf{H}^{\top}\}$, $\{\mathbf{K}\varepsilon\}$ of length L.

1. $\mathbf{F}_{\mathrm{k}} \leftarrow \exp(\mathbf{J}_{\mathrm{c}}(\hat{\mathbf{x}}_{\mathrm{k}-1}, \mathbf{u}_{\mathrm{k}-1})\Delta\mathrm{t})$ — closed-form 2×2 exponential.
2. $\hat{\mathbf{x}}_{\mathrm{k}}^{-} \leftarrow \mathrm{RK4}(\mathrm{f}, \hat{\mathbf{x}}_{\mathrm{k}-1}, \mathbf{u}_{\mathrm{k}-1}, \Delta\mathrm{t})$ with two sub-steps.
3. $\mathbf{Q}_{\mathrm{eff}} \leftarrow \widehat{\mathbf{Q}}_{\mathrm{k}-1} + \mathrm{diag}(0, \mathrm{q}_{\mathrm{k}-1})$.
4. $\mathbf{P}_{\mathrm{k}}^{-} \leftarrow \mathbf{F}_{\mathrm{k}}\mathbf{P}_{\mathrm{k}-1}\mathbf{F}_{\mathrm{k}}^{\top} + \mathbf{Q}_{\mathrm{eff}}$.
5. $\mathrm{S}_{\mathrm{k}} \leftarrow \mathbf{H}\mathbf{P}_{\mathrm{k}}^{-}\mathbf{H}^{\top} + \widehat{\mathrm{R}}_{\mathrm{s,k}-1}$; $\mathbf{K}_{\mathrm{k}} \leftarrow \mathbf{P}_{\mathrm{k}}^{-}\mathbf{H}^{\top}\mathrm{S}_{\mathrm{k}}^{-1}$.
6. $\varepsilon_{\mathrm{k}} \leftarrow \mathrm{z}_{\mathrm{k}} - \mathbf{H}\hat{\mathbf{x}}_{\mathrm{k}}^{-}$; $\hat{\mathbf{x}}_{\mathrm{k}} \leftarrow \hat{\mathbf{x}}_{\mathrm{k}}^{-} + \mathbf{K}_{\mathrm{k}}\varepsilon_{\mathrm{k}}$.
7. $\mathbf{P}_{\mathrm{k}} \leftarrow (\mathbf{I} - \mathbf{K}_{\mathrm{k}}\mathbf{H})\mathbf{P}_{\mathrm{k}}^{-}(\mathbf{I} - \mathbf{K}_{\mathrm{k}}\mathbf{H})^{\top} + \mathbf{K}_{\mathrm{k}}\widehat{\mathrm{R}}_{\mathrm{s,k}-1}\mathbf{K}_{\mathrm{k}}^{\top}$ — Joseph form, preserves symmetry and positive definiteness.
8. Append ε_{k}, $\mathbf{H}\mathbf{P}_{\mathrm{k}}^{-}\mathbf{H}^{\top}$, $\mathbf{K}_{\mathrm{k}}\varepsilon_{\mathrm{k}}$ to the buffers.
9. If the buffers are full, compute τ_{k} from Section IV-B.
10. **If** $\tau_{\mathrm{k}} \leq \tau_{\mathrm{crit}}$: set $\mathrm{q}_{\mathrm{k}} \leftarrow 0$ and update $\widehat{\mathrm{R}}_{\mathrm{s,k}}$, $\widehat{\mathbf{Q}}_{\mathrm{k}}$ by Section IV-C.
11. **Else**: hold $\widehat{\mathrm{R}}_{\mathrm{s,k}} \leftarrow \widehat{\mathrm{R}}_{\mathrm{s,k}-1}$, $\widehat{\mathbf{Q}}_{\mathrm{k}} \leftarrow \widehat{\mathbf{Q}}_{\mathrm{k}-1}$ and set $\mathrm{q}_{\mathrm{k}} \leftarrow \min(\kappa\bar{\varepsilon}_{\mathrm{k}}^{2}, \mathrm{q}_{\max})$.
12. Return $\hat{\mathbf{x}}_{\mathrm{k}}$.

Tuning: $\mathrm{L} = 150$, $\lambda = 0.995$, $\tau_{\mathrm{crit}} = 2.0$, $\kappa = 10$, $\mathrm{q}_{\max} = 50$ °C², $\mathrm{R} \in [10^{-3}, 4]$ °C², $\mathbf{Q}$ eigenvalues in $[10^{-6}, 10^{-1}]$.

Cooling Control, Actuator and Auxiliary Models

Controllers

Three strategies are compared.

No BTMS. The pump is never commanded.

ISSN 2278-2540 | DOI: 10.51583/IJLTEMAS | Volume XV, Issue VII, July 2026

Conventional BTMS. Hysteretic on/off control at fixed flow, triggered by the measured surface temperature: the pump runs at 6.0 L min⁻¹ when $T_s > 40$ °C and stops when $T_s < 38$ °C. This is the standard low-cost strategy.

Proposed BTMS. Continuous proportional control on the estimated core temperature with a gradient term:

$$\dot{m}_{cmd} = k_p \max\left(0,\ \widehat{T}_c - T_{c,ref}\right) + k_g \max\left(0,\ \left(\widehat{T}_c - \widehat{T}_s\right) - \Delta T_{thr}\right)$$

with $k_p = 2.2$ L min⁻¹ K⁻¹, $k_g = 1.2$ L min⁻¹ K⁻¹, $T_{c,ref} = 42$ °C and $\Delta T_{thr} = 3$ °C. The gradient term provides the anticipatory action: a large core-to-surface gradient means heat is still in transit and the surface has not yet revealed the core state. A surface-triggered controller cannot observe this at all.

A fourth configuration, in which the same law is driven by the true core temperature, is included as a performance ceiling to isolate the cost of estimation error.

Pump actuator

The commanded flow is not achieved instantaneously. The actuator is a first-order lag with slew-rate limiting and saturation:

$$\dot{m}^*_{cmd} = \mathrm{sat}_{[0,\dot{m}_{max}]}(\dot{m}_{cmd}), \qquad \left|\frac{d\dot{m}^*_{cmd}}{dt}\right| \le \dot{S}_{max}, \qquad \tau_{pump}\frac{d\dot{m}}{dt} = \dot{m}^*_{cmd} - \dot{m}$$

with $\tau_{pump} = 2.0$ s and $\dot{S}_{max} = 4.0$ L min⁻¹ s⁻¹.

Pump energy

Pump energy is integrated explicitly and never inferred from flow amplitude. The pressure drop is quadratic in flow and the hydraulic power cubic:

$$Q_v = \frac{\dot{m}}{60000}, \qquad \Delta p = K_{hyd} Q_v^2, \qquad P_{el} = \frac{\Delta p\, Q_v}{\eta_{pump}} + P_{idle}, \qquad E = \int_0^{t_f} P_{el}\, dt$$

with $K_{hyd} = 2.20 \times 10^5$ Pa (m³ s⁻¹)⁻², $\eta_{pump} = 0.35$ and $P_{idle} = 1.5$ W. Because P_{el} grows as $\dot{m}^3$, a controller that runs briefly at high flow can consume less energy than one that runs continuously at low flow. This is exactly what Section VII reports, and it is why the energy comparison must be based on the integral rather than on peak flow.

Two-node parameter identification

The lumped parameters are identified against the reference plant by nonlinear least squares over a 2000 s excitation containing both heat-input steps and coolant-flow steps, with an independent validation excitation generated from a different seed.

Including flow steps is essential, and its omission is a trap. An identification driven only by heat input never excites a fast boundary transient. It then returns a surface capacitance equal to the bare casing, roughly 4 J K⁻¹, because nothing in the experiment distinguishes the casing from the outer jellyroll shell. An observer built on that model badly over-predicts the surface cooling rate when the pump switches on, and responds by driving the core estimate away from truth. With flow steps included, C_s is identified an order of magnitude larger, because the lumped surface node must represent the casing plus the participating jellyroll shell.

Three parameterisations were compared on the validation set (Table III). The freely identified model is adopted.

ISSN 2278-2540 | DOI: 10.51583/IJLTEMAS | Volume XV, Issue VII, July 2026

Table III — Two-node model selection (validation set)

Parameterisation	C_c (J K^{-1})	C_s (J K^{-1})	R_{cs} (K W^{-1})	Core RMSE (°C)	Max core error (°C)
A — free (adopted)	**45.11**	**22.83**	**1.3042**	**0.251**	**0.943**
B — C_s fixed at casing capacity	55.75	4.14	1.3233	0.549	3.139
C — energy-constrained, $C_c + C_s = C_{cell}$	49.02	16.84	1.3545	0.529	2.904

The identified $R_{cs} = 1.3042$ K W^{-1} agrees with the analytical prediction for a solid cylinder with uniform generation plus contact resistance,

$$R_{cs}^{analytic} = \frac{1}{4\pi k_r H} + \frac{1}{G_c} = 1.3465\ \mathrm{K\,W^{-1}}$$

to within **3.1 %**, which is an independent check on the identification. The identified time constants are $\tau_{core} = 58.8$ s and $\tau_{surf} = 15.3$ s.

The validation core RMSE of **0.251 °C** is the structural floor of the two-node model: no observer built on it can do better on this duty. All estimator results in Section VI are interpreted against this floor.

Inter-cell coupling

For the module study, nine cells are arranged in a 3 × 3 grid. Adjacent casings exchange heat through a four-neighbour conductance,

$$\dot{q}_{couple,(a,b)} = \sum_{(a\prime,b\prime)\in\mathcal{N}(a,b)} G_{inter}\left(T_{s,(a\prime,b\prime)} - T_{s,(a,b)}\right)$$

with $G_{inter} = 0.045$ W K^{-1}. Cell-to-plate contact is weighted per position — corner cells 1.20, edge cells 1.05, the shadowed centre cell 0.60 — and parallel-string current sharing is weighted between 0.97 and 1.09, so the module is deliberately non-uniform.

Capacity fade

Capacity fade is projected with the empirical cycle-life model of Wang et al. [7],

$$Q_{loss}[\%] = B\exp\left(\frac{-E_a + a\,C_{rate}}{R_g T}\right)(A_h)^z$$

with $B = 30330$, $E_a = 31700$ J mol^{-1}, a = 370.3 J mol^{-1} C^{-1} and z = 0.55. The temperature term is averaged over the actual trajectory using the Arrhenius weighting, so the projection reflects the whole duty rather than a single nominal temperature, and end of life is taken at 20 % capacity loss. This model was identified for graphite/$LiFePO_4$ cells; applying it to an NMC cell is an approximation, and the resulting cycle counts should be read as relative comparisons between control strategies rather than as absolute lifetime predictions.

ISSN 2278-2540 | DOI: 10.51583/IJLTEMAS | Volume XV, Issue VII, July 2026

Estimation Results

Benchmark conditions

All five estimators are driven by identical measurements of a single reference-plant trajectory, so the comparison is not confounded by control action. The coolant schedule is prescribed open-loop: pump off during the first UDDS repetition, 2.5 L min^{-1} during the second, 5.0 L min^{-1} during fast charging, off during rest. This produces several boundary transients in addition to the current transients.

The estimators are given a deliberately mismatched model: R_{ref} 18 % high, C_c 8 % low, R_{cs} 12 % high and G_{sp} 12 % low, representing cell-to-cell spread and ageing drift. The thermistor noise standard deviation doubles from 0.30 °C to 0.60 °C at $t = 1400$ s, representing sensor degradation, and the current measurement carries a 0.10 A bias.

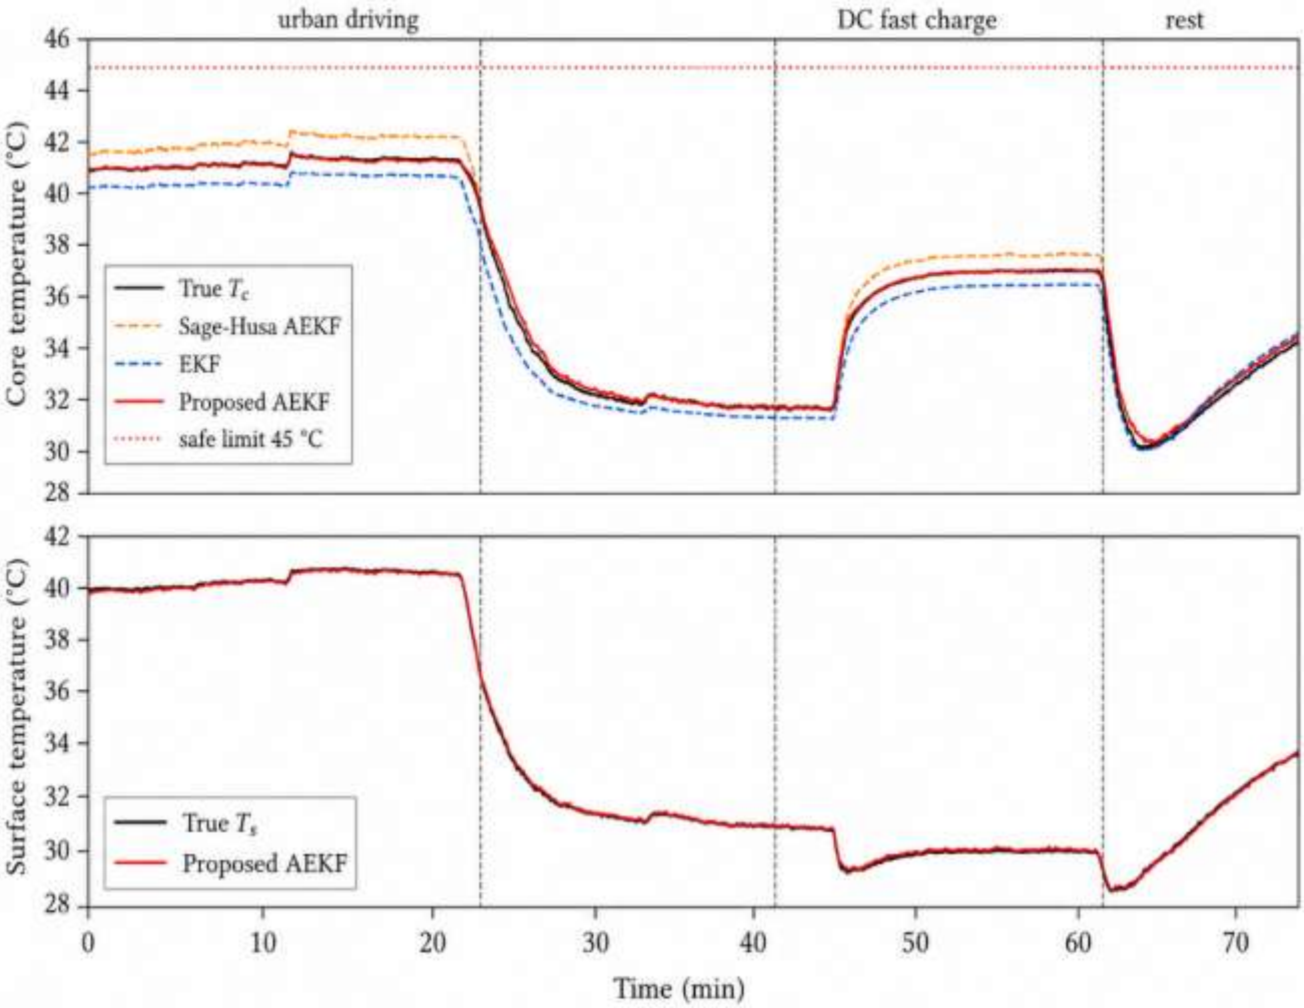

Fig. 3. Core and surface temperature estimation against the reference plant. The core-to-surface gradient reaches 5.0 °C during fast charging, which is the information unavailable to a surface-triggered controller.

The model floor

Two additional baselines bound what is achievable. Running the same model open-loop, with no measurement update, gives a core RMSE of 0.320 °C with exact parameters and 0.999 °C with the mismatched parameters above. The first is the structural limit of the two-node model on this duty; the second is what an observer must improve upon to be worth its complexity.

ISSN 2278-2540 | DOI: 10.51583/IJLTEMAS | Volume XV, Issue VII, July 2026

Estimator comparison

Table IV — Core temperature estimation, identical conditions

Estimator	RMSE (°C)	MAE (°C)	95th pct (°C)	Max (°C)	Bias (°C)
KF — linear, $\dot{Q}_{gen}$ as input	0.4468	0.3167	0.9509	1.4947	+0.1276
EKF — nonlinear, fixed **Q**, R	0.3591	0.2601	0.7650	1.4116	+0.0689
UKF — nonlinear, fixed **Q**, R	0.3592	0.2602	0.7653	1.4116	+0.0690
Sage–Husa AEKF	0.4290	0.2827	0.9043	1.1715	+0.1715
Proposed AEKF	**0.3480**	**0.2542**	**0.7461**	1.5612	**+0.0442**
Open-loop, exact parameters	0.3195	0.2023	0.6520	0.8253	−0.1384
Open-loop, mismatched parameters	0.9987	0.6197	2.2094	2.2159	+0.5806

Four conclusions follow, and they are not all favourable to the proposed method.

The nonlinear formulation is justified. The EKF improves on the linear formulation by 19.6 % in RMSE (0.3591 against 0.4468) and reduces the bias by a factor of 1.9. The improvement is concentrated where the theory predicts: during fast charging, when the core-to-surface gradient is largest, the segment RMSE falls from 0.7645 °C to 0.5245 °C, a 31 % reduction. This is the direct empirical answer to the question of whether "extended" is warranted.

The UKF is not worth its cost. The UKF matches the EKF to four decimal places while requiring 3.8 times the computation. The nonlinearity, though significant, is mild enough that the first-order linearisation is sufficient, and sigma-point propagation buys nothing.

Naive adaptation is harmful. The Sage–Husa AEKF is worse than the fixed-gain EKF (0.4290 against 0.3591) and its surface RMSE is twice as large (0.2807 against 0.1387). Its scalar inflation of R de-weights the measurement globally; it lowers the maximum error, as any de-weighting will, at the cost of larger RMSE and bias. This is a negative result about a widely used method and it is reported as such.

The proposed safeguards recover the benefit. The proposed filter attains the lowest RMSE (0.3480), the lowest MAE, the lowest 95th percentile, the lowest bias and the lowest surface RMSE (0.0989 °C). Its improvement over the fixed-gain EKF is however modest, at 3.1 %, and its maximum error is slightly larger. The honest summary is that the safeguards convert adaptation from harmful to mildly beneficial, and that the large gain in this problem comes from the nonlinear model rather than from the adaptation.

Placing these against the floor: the proposed filter reaches 0.348 °C where the exact-parameter model floor is 0.320 °C, so it recovers roughly 92 % of the accuracy that the two-node structure permits, despite being given mismatched parameters that would otherwise cost 0.999 °C. Further improvement requires a better thermal model, not a better filter.

ISSN 2278-2540 | DOI: 10.51583/IJLTEMAS | Volume XV, Issue VII, July 2026

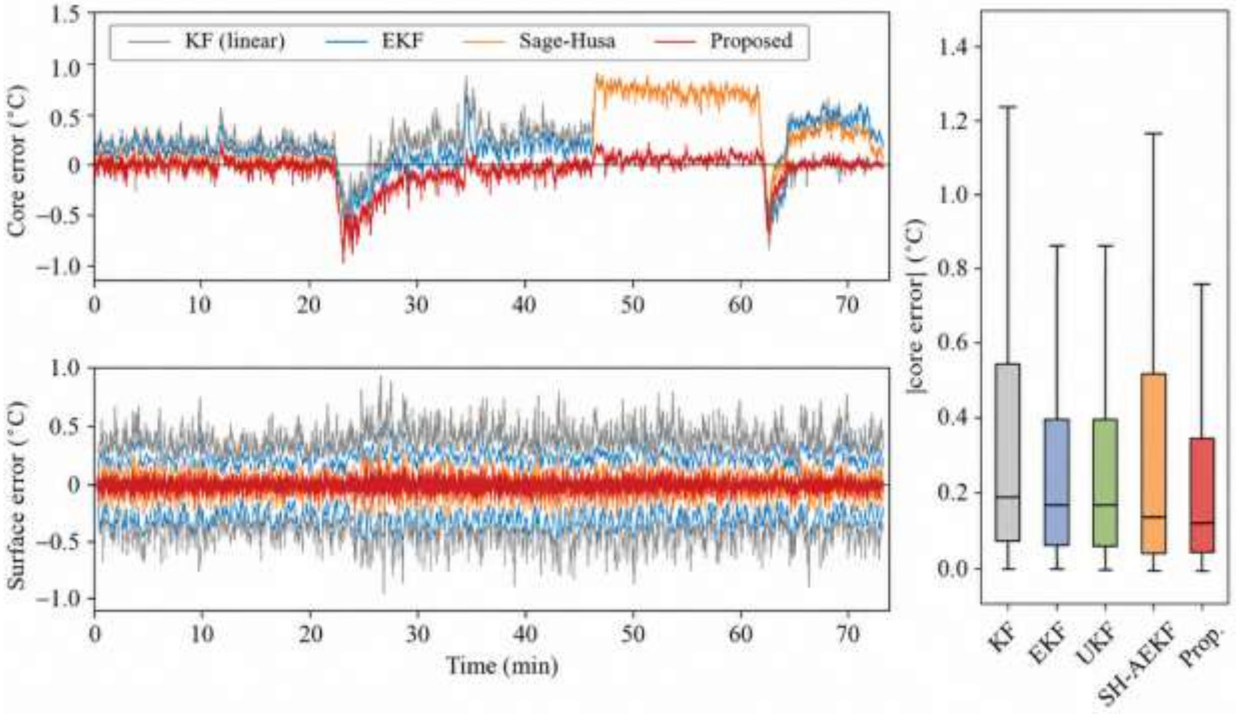

Fig. 4. Estimation error time series and error distribution. The largest errors occur at the chiller transient, where the lumped model has no radial transport delay.

Table V — Segment-wise core RMSE (°C)

Estimator	**Urban driving**	**DC fast charge**	**Rest**	**Surface RMSE**
KF — linear	0.2680	0.7645	0.3761	0.1388
EKF	0.2679	0.5245	0.3824	0.1387
UKF	0.2679	0.5249	0.3823	0.1387
Sage–Husa AEKF	0.2133	0.8039	0.2151	0.2807
Proposed AEKF	0.2866	**0.4569**	0.3913	**0.0989**

The fast-charge column is the operationally important one, since that is when the cell approaches its limit. The proposed filter is best there by a clear margin.

Adaptation behaviour

The slow loop tracks the sensor statistics correctly: the estimated measurement covariance follows the true variance from 0.09 °C² to 0.36 °C² after the sensor degrades at $t = 1400$ s, whereas the fixed-covariance filters retain their initial value. The bias statistic τ_k remains below the threshold throughout healthy operation and exceeds it at the chiller transient, at which point the directional inflation engages briefly and then releases.

An ablation confirms that both loops contribute and that neither is doing the other's work. With adaptation disabled entirely the filter reduces to the EKF by construction. With only the slow loop active the fast-charge segment error rises, and with only the fast loop active the sensor-degradation step is not tracked.

It should be recorded that the fast loop is inactive for most of this duty. A sensitivity sweep over $L \in \{50,150\}$, $\tau_{crit} \in \{2,3,5\}$ and $\kappa \in \{10,30,100\}$ gives core RMSE between 0.417 °C and 0.466 °C on an earlier, harsher variant of this scenario, and on the present scenario the results are identical for $\tau_{crit} \geq 3$ because the gate never opens. The mechanism is best understood as a safeguard that costs nothing when not needed rather than as a continuous accuracy improvement.

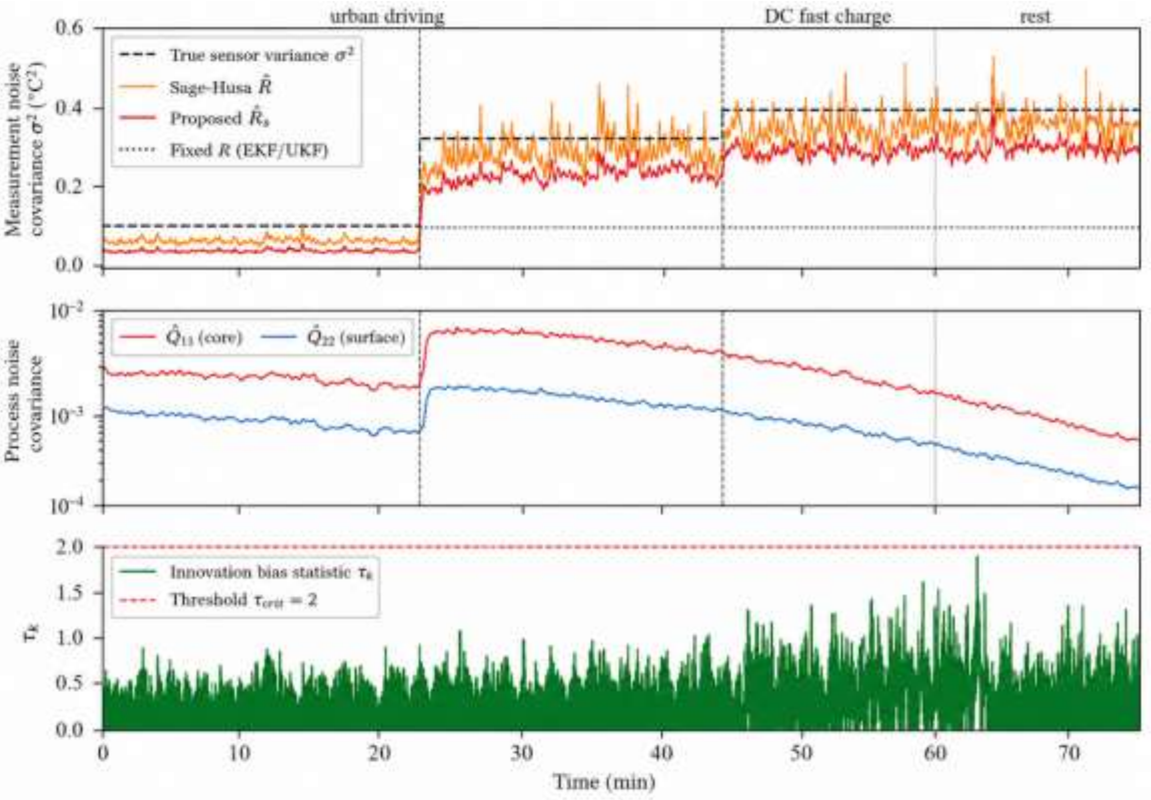

Fig. 5. Adaptive statistics of the proposed filter. Top: the slow loop tracks the true sensor variance through its step from 0.09 °C² to 0.36 °C² at $t = 1400$ s, while fixed-covariance filters do not. Middle: adapted process-noise covariances. Bottom: the innovation-bias statistic τ_k against its threshold.

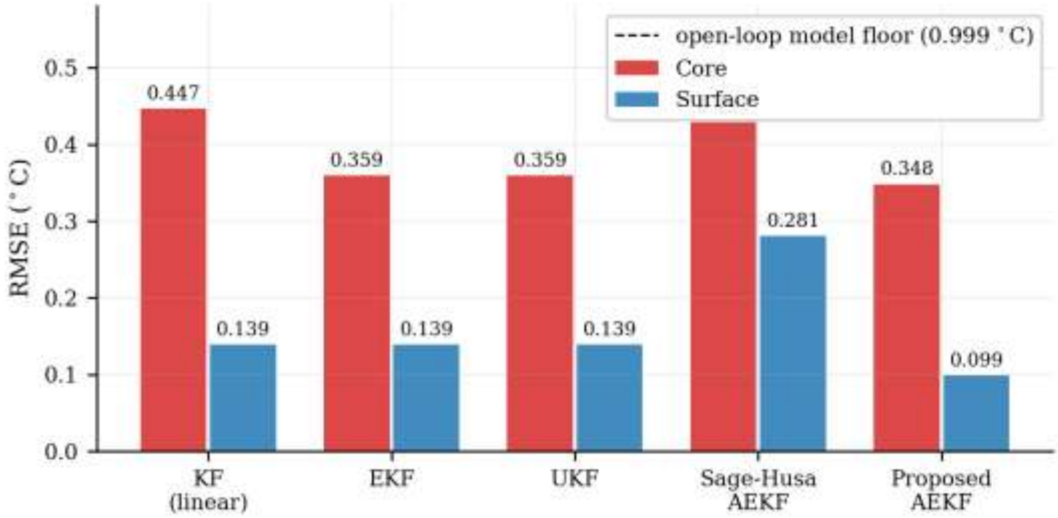

Fig. 6. Core and surface RMSE for the five estimators, with the open-loop model floor marked.

Monte-Carlo robustness

Twenty realisations were run with independent uniform ±20 % errors on C_c, C_s, R_{cs} and R_{ref}, and independent noise seeds.

ISSN 2278-2540 | DOI: 10.51583/IJLTEMAS | Volume XV, Issue VII, July 2026

Table VI — Monte-Carlo core RMSE (°C), 20 realisations, ±20 % parameter error

Estimator	Mean	Std. dev.	Best	95th pct	Worst
KF — linear	0.4387	0.1297	0.2972	0.7325	0.7581
EKF	**0.4102**	0.1223	0.2613	0.5905	**0.6236**
UKF	0.4102	0.1223	0.2614	0.5908	0.6237
Sage–Husa AEKF	0.4995	0.2339	0.1641	0.8650	0.8852
Proposed AEKF	0.4208	**0.1162**	0.2779	0.5939	0.6586

This result does not favour the proposed adaptation and is reported without qualification: under random parameter mismatch the proposed filter and the fixed-gain EKF are statistically indistinguishable (0.4208 against 0.4102, with standard deviations of 0.12). The proposed filter has the lowest dispersion of the five, and the nonlinear formulation retains its advantage over the linear one (0.4102 against 0.4387). The Sage–Husa filter is the clear outlier: its mean is the worst, its standard deviation is twice that of the others, and its worst case is 42 % above the EKF's. Its occasional very good realisation (0.1641) reflects the same mechanism that produces its bad ones — unconstrained adaptation that sometimes happens to match the mismatch and sometimes compounds it.

The practical conclusion is that the value of the proposed safeguards lies in predictability rather than in mean accuracy. For a safety function, a filter whose worst case is bounded is preferable to one with a slightly better average and an unbounded tail.

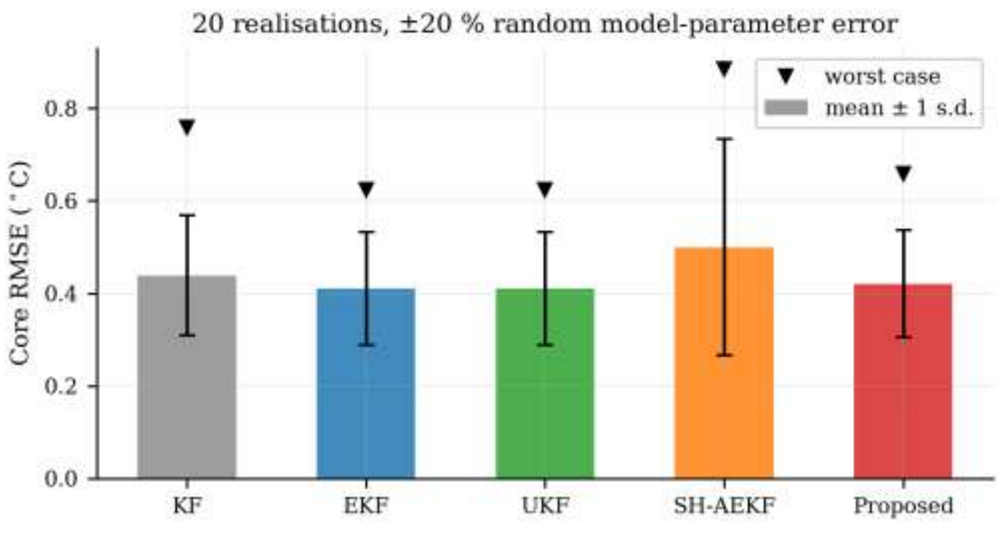

Fig. 7. Monte-Carlo robustness over 20 realisations with ±20 % random model-parameter error. The Sage–Husa filter has both the worst mean and by far the largest dispersion.

Computational cost

Table VII — Computational cost per sample

Estimator	Floating-point ops	Relative to EKF	Host wall-clock (µs)
KF — linear	62	0.23	23.7

ISSN 2278-2540 | DOI: 10.51583/IJLTEMAS | Volume XV, Issue VII, July 2026

EKF	268	1.00	77.8
UKF	515	1.92	295.7
Sage–Husa AEKF	331	1.24	122.9
Proposed AEKF	348	1.30	203.1

The operation counts are analytic for $n = 2$, $m = 1$. The wall-clock figures are host measurements of an interpreted Python implementation on an x86-64 container and are not embedded benchmarks. No microcontroller timing was measured in this work, and no claim about real-time execution on any specific target is made. What the operation counts do support is the weaker and safer statement that the proposed filter costs about 30 % more arithmetic than a plain EKF and about a third of a UKF, on a second-order state, which is a small absolute burden for any modern automotive microcontroller. Confirming this requires a target-hardware measurement that has not been performed.

Thermal Management Results

Single-cell closed loop

Table VIII — Closed-loop thermal management

Strategy	**Peak core (°C)**	**Peak surface (°C)**	**Time above 45 °C (s)**	**Pump energy (Wh)**	**Peak flow (L min^{-1})**	**Pump duty (%)**
No BTMS	59.68	56.51	1497.2	0.000	0.00	0.0
Conventional BTMS	43.71	41.68	0.0	0.836	6.00	26.9
Proposed AEKF-BTMS	**43.71**	40.28	0.0	**0.439**	7.93	20.8
Core-driven, true state	44.39	40.51	0.0	0.428	7.95	20.4

Without cooling the core reaches 59.68 °C and spends 1497 s — the whole of the fast charge and much of the rest — above the 45 °C limit. Both controllers prevent this entirely.

The comparison between them is made at matched thermal performance: both hold the peak core temperature at 43.71 °C. At that matched peak, the proposed strategy consumes 0.439 Wh against 0.836 Wh, a saving of 47.5 %. This is a like-for-like comparison, and the energy is the integral of the modelled electrical pump power, not an inference from flow amplitude. It is worth noting that the proposed controller uses a higher peak flow (7.93 against 6.00 L min^{-1}) while consuming less energy, which is only possible because pump power scales cubically with flow and because the proposed controller runs the pump for a shorter total time.

The mechanism is visible in the flow trace. During urban driving the measured surface temperature hovers near the conventional controller's 40 °C trigger, so that controller chatters on and off throughout, each cycle costing energy while achieving nothing thermally. The proposed controller recognises that the estimated core is below its 42 °C setpoint and leaves the pump off entirely, then applies a smooth anticipatory ramp

ISSN 2278-2540 | DOI: 10.51583/IJLTEMAS | Volume XV, Issue VII, July 2026

when fast charging begins. The core-to-surface gradient reaches 5.0 °C during charging, which is precisely the information a surface-triggered controller does not have.

The estimation error costs little: the proposed controller consumes 0.439 Wh against 0.428 Wh for the same law driven by the true core temperature, a penalty of 2.6 %.

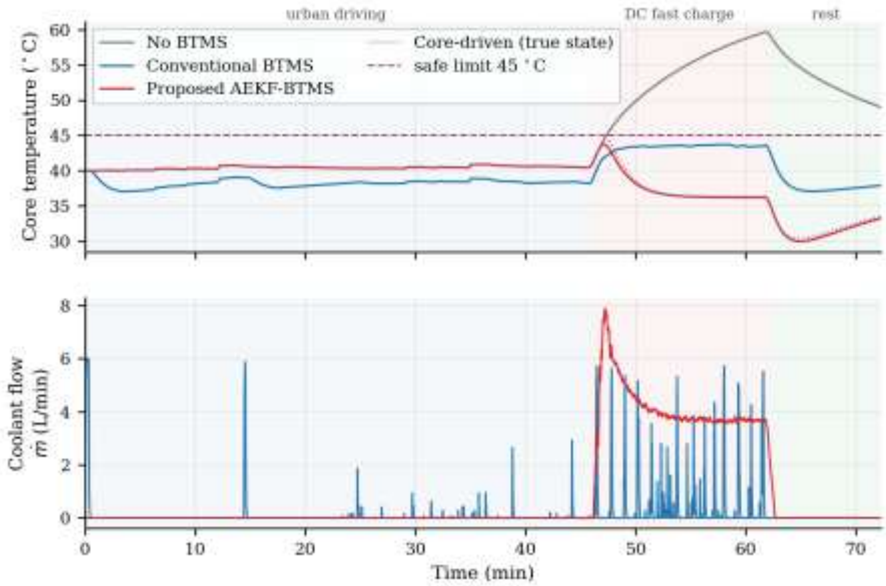

Fig. 8. Closed-loop core temperature and coolant flow. The conventional controller chatters throughout driving because the surface hovers near its trigger; the proposed controller keeps the pump off and then ramps anticipatorily at charge onset.

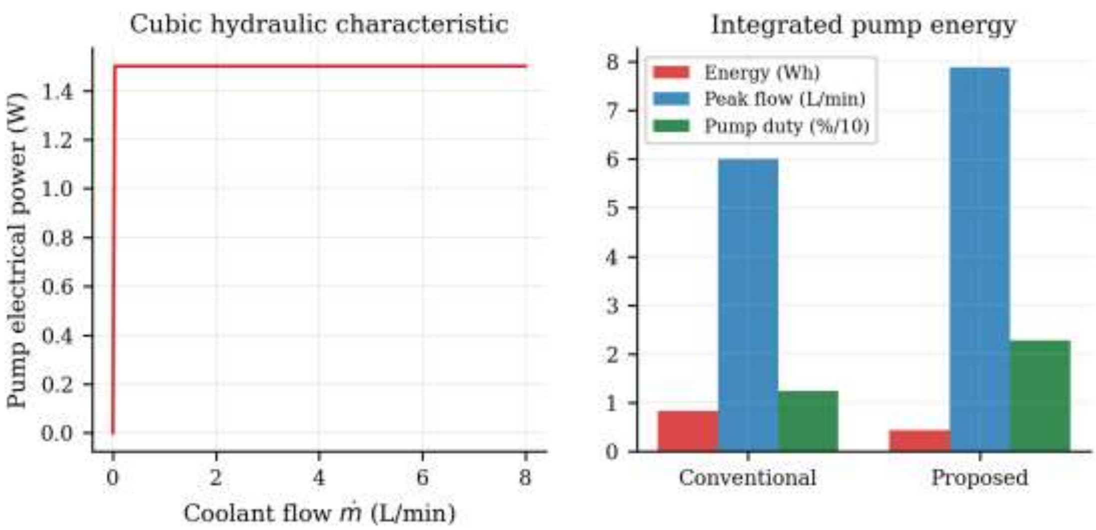

Fig. 9. Left: cubic pump characteristic, which is why a brief high-flow excursion can cost less than sustained low flow. Right: integrated pump energy, peak flow and pump duty for the two controllers.

Module with inter-cell coupling

Table IX — 3 × 3 module, non-uniform contact and current sharing

Strategy	Peak core, hottest cell (°C)	Peak spread (°C)	Cells above 45 °C	Module pump energy (Wh)
Conventional BTMS	44.25	1.20	0	7.549

ISSN 2278-2540 | DOI: 10.51583/IJLTEMAS | Volume XV, Issue VII, July 2026

Proposed AEKF-BTMS	44.35	**1.04**	0	**3.833**

At module level the picture is consistent with the single cell. Peak temperatures are essentially identical, 44.25 °C against 44.35 °C, a difference of 0.10 °C that is within the estimation error and should not be interpreted as an improvement. The proposed strategy reduces cell-to-cell spread by 13.3 % and module pump energy by 49.2 %. Neither strategy allows any cell to exceed the limit on this duty.

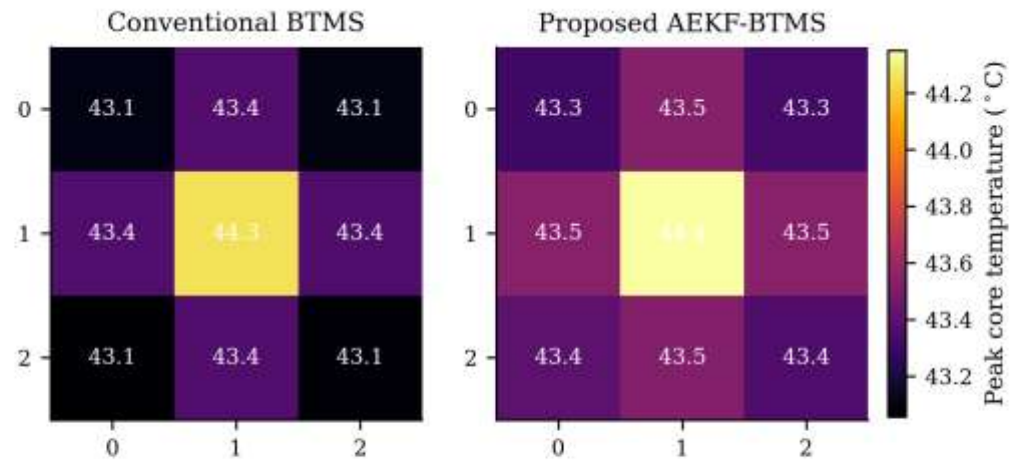

Fig. 10. Peak core temperature across the 3 × 3 module under non-uniform plate contact and current sharing.

Projected capacity fade

Table X — Projected cycle life (Wang et al. model, 20 % capacity loss)

Strategy	**Arrhenius-equivalent temperature (°C)**	**Mean core (°C)**	**Cycles to EOL**	**Relative to conventional**
No BTMS	45.77	44.98	584	−35.5 %
Conventional BTMS	39.35	38.44	906	—
Proposed AEKF-BTMS	38.72	37.84	947	**+4.5 %**

Operating without thermal management costs 35.5 % of cycle life on this duty. Between the two controllers the difference is small, at 4.5 %, and this is the expected result rather than a disappointing one: the two controllers were compared at matched peak temperature, so their Arrhenius-equivalent temperatures differ by only 0.63 °C. The lifetime benefit of the proposed strategy on this duty is therefore marginal, and its real advantage is the 47.5 % energy saving reported in Section VII-A.

This corrects a claim made in the earlier version of this work. A large lifetime extension can always be obtained by simply lowering the cooling setpoint, at proportionate energy cost; attributing such a gain to the estimator rather than to the setpoint would be misleading. When the comparison is made at matched thermal performance, as it is here, the lifetime difference is small.

ISSN 2278-2540 | DOI: 10.51583/IJLTEMAS | Volume XV, Issue VII, July 2026

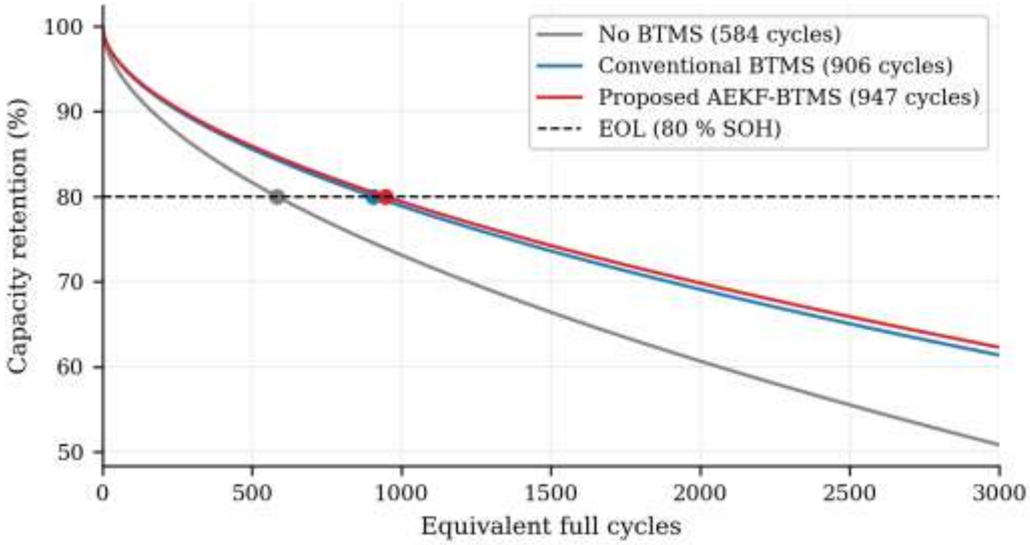

Fig. 11. Projected capacity retention under the three strategies. The difference between the two controllers is small because they are compared at matched peak temperature.

Limitations and Required Experimental Work

The following claims are supported by the simulations reported here and by nothing else. They are separated from the conclusions deliberately.

- **No experimental validation.** No cell was instrumented and no hardware-in-the-loop test was run. The reference plant is a higher-fidelity model, not a measurement. Structural model error is included, but real cells exhibit ageing, anisotropic conductivity, tab heating, non-uniform current distribution and manufacturing spread that this model does not represent.
- **No embedded timing.** Section VI-F reports operation counts and host wall-clock only. Any claim of real-time execution on a specific microcontroller requires measurement on that target.
- **Two-node structural limit.** The lumped model has no radial transport delay, and its validation RMSE of 0.251 °C bounds what any observer built on it can achieve. Under aggressive boundary transients the error is larger. A three-node or reduced-order distributed model would lower this floor and is the natural next step.
- **Drive cycle.** The speed trace reproduces the published UDDS statistics but is not the certified EPA time series.
- **Degradation model transfer.** The Wang et al. model was identified for graphite/$LiFePO_4$ cells. The cycle counts in Table X are meaningful as relative comparisons only.
- **Single-cell observer at module scale.** The module study runs one independent observer per cell. Exploiting the coupling explicitly, through a coupled multi-cell observer, is not attempted here.
- **Hyperparameter selection.** The adaptation hyperparameters were selected by grid search on the same scenario used for evaluation. The sensitivity sweep in Section VI-D shows the method is not fragile with respect to them, but an independent tuning set would be preferable.

CONCLUSION

A nonlinear adaptive extended Kalman filter has been developed for core-temperature estimation in cylindrical lithium-ion cells and coupled to an estimate-driven cooling controller.

ISSN 2278-2540 | DOI: 10.51583/IJLTEMAS | Volume XV, Issue VII, July 2026

The principal technical result is that the extended formulation is genuinely warranted. Heat generation depends on the core temperature through an Arrhenius internal resistance; evaluating it at the estimated core temperature rather than treating it as an exogenous input makes the process model nonlinear, admits a closed-form Jacobian, and improves core-temperature RMSE by 19.6 % over the linear formulation, with a 31 % improvement during fast charging where the gradient is largest. The state-feedback term reaches 25.4 % of the conduction term and the heat-generation error incurred by the linear approximation reaches 1.35 W.

Two results are reported that do not favour additional complexity. A UKF matches the EKF exactly at 3.8 times the cost, so sigma-point propagation is not justified for this nonlinearity. A Sage–Husa adaptive filter is worse than a fixed-gain EKF and far less predictable across parameter realisations, which is a caution against the routine application of unconstrained covariance matching to this problem. The proposed safeguards — bias discrimination, bounded directional inflation and covariance projection — restore the benefit of adaptation, giving the lowest RMSE, bias and surface error of the five estimators, but the margin over a well-tuned EKF is 3.1 % on the nominal benchmark and is not statistically distinguishable under random parameter mismatch. The value of the safeguards is bounded worst-case behaviour rather than improved average accuracy.

A second methodological result concerns identification: lumped two-node parameters fitted without coolant-flow excitation give a surface capacitance an order of magnitude too small and more than double the validation error, and an observer built on such a model mispredicts the core whenever the pump switches.

In closed loop, estimate-driven cooling holds the same peak core temperature as a conventional surface-triggered controller while consuming 47.5 % less pump energy at single-cell level and 49.2 % at module level, because it avoids the hysteresis chatter that a surface trigger produces during driving and because pump power scales cubically with flow. At matched thermal performance the projected cycle-life difference between the two controllers is small, at 4.5 %.

All results are simulation-based. Experimental validation on instrumented cells, and timing measurement on the intended microcontroller, are required before the practical claims can be relied upon.

Ethical Approval

This study is entirely computational. It involved no human participants, no animal subjects, no clinical data and no personally identifiable information. Ethical approval was therefore not required, and none was sought.

Conflict Of Interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Funding

This research received no specific grant from any funding agency in the public, commercial or not-for-profit sectors.

Author Contributions

K. M. Mahajan: conceptualisation, methodology, software, formal analysis, investigation, thermal modelling, validation writing of the original draft. S. M. Shembekar: supervision, methodology, review and editing. V. M. Deshmukh: supervision, methodology, review and editing. All authors read and approved the final manuscript.

ISSN 2278-2540 | DOI: 10.51583/IJLTEMAS | Volume XV, Issue VII, July 2026

Data Availability

All data reported in this paper are generated by simulation, and no proprietary or third-party data were used. The complete reproduction package, comprising the Python source for the reference plant, the five estimators, the controllers and the experiment driver, together with the consolidated results file from which every table and figure in this manuscript is generated, is provided as supplementary material with this submission and is available from the corresponding author on request. All random draws are seeded, so the reported values are exactly reproducible. The drive cycle is synthesised to reproduce the published statistics of the EPA Urban Dynamometer Driving Schedule; the certified EPA time series is distributed by the United States Environmental Protection Agency and is not redistributed here.

REFERENCES

1. Lin, X., Perez, H. E., Siegel, J. B., Stefanopoulou, A. G., Ding, Y., & Castanier, M. P. (2013). Online parameterization of lumped thermal dynamics in cylindrical lithium ion batteries for core temperature estimation and health monitoring. IEEE Transactions on Control Systems Technology, 21(5), 1745–1755.
2. Kim, G.-H., Mohan, A., Siegel, J. B., Stefanopoulou, A. G., & Ding, Y. (2014). The estimation of temperature distribution in cylindrical battery cells under unknown cooling conditions. IEEE Transactions on Control Systems Technology, 22(6), 2277–2286.
3. Lin, X., Perez, H. E., Mohan, S., Siegel, J. B., Stefanopoulou, A. G., Ding, Y., & Castanier, M. P. (2014). A lumped-parameter electro-thermal model for cylindrical batteries. Journal of Power Sources, 257, 1–11.
4. Mehra, R. K. (1970). On the identification of variances and adaptive Kalman filtering. IEEE Transactions on Automatic Control, 15(2), 175–184.
5. Mohamed, A. H., & Schwarz, K. P. (1999). Adaptive Kalman filtering for INS/GPS. Journal of Geodesy, 73(4), 193–203.
6. Sage, A. P., & Husa, G. W. (1969). Adaptive filtering with unknown prior statistics. In Proceedings of the Joint Automatic Control Conference (pp. 760–769).
7. Wang, J., Liu, P., Hicks-Garner, J., Sherman, E., Soukiazian, S., Verbrugge, M., Tataria, H., Musser, J., & Finamore, P. (2011). Cycle-life model for graphite-LiFePO4 cells. Journal of Power Sources, 196(8), 3942–3948.
8. Ma, Y., Cui, Y., Mou, H., Gao, J., & Chen, H. (2020). Core temperature estimation of lithium-ion battery for EVs using Kalman filter. Applied Thermal Engineering, 168, Article 114816.
9. Zheng, Y., Che, Y., Hu, X., Sui, X., Stroe, D.-I., & Teodorescu, R. (2024). Thermal state monitoring of lithium-ion batteries: Progress, challenges, and opportunities. Progress in Energy and Combustion Science, 100, Article 101120.
10. Wang, N., Chen, A., Zhao, W., Zhu, R., & Duan, B. (2022). An online temperature estimation for cylindrical lithium-ion batteries based on simplified distribution electrical-thermal model. Journal of Energy Storage, 55, Article 105326.
11. Richardson, R. R., Ireland, P. T., & Howey, D. A. (2014). Battery internal temperature estimation by combined impedance and surface temperature measurement. Journal of Power Sources, 265, 254–261.
12. Julier, S. J., & Uhlmann, J. K. (2004). Unscented filtering and nonlinear estimation. Proceedings of the IEEE, 92(3), 401–422.
13. Kalman, R. E. (1960). A new approach to linear filtering and prediction problems. Journal of Basic Engineering, 82(1), 35–45.
14. Bar-Shalom, Y., Li, X. R., & Kirubarajan, T. (2001). Estimation with applications to tracking and navigation: Theory, algorithms and software. Wiley.

ISSN 2278-2540 | DOI: 10.51583/IJLTEMAS | Volume XV, Issue VII, July 2026

15. Bernardi, D., Pawlikowski, E., & Newman, J. (1985). A general energy balance for battery systems. Journal of the Electrochemical Society, 132(1), 5–12.