Article | Received 21 September 2026; Revised 9 November 2026; Accepted 16 November 2026; Published 27 November 2026
https://doi.org/10.55092/mt20260009

Uncertainty-aware discrete-time barrier filtering for hybrid energy storage under unpreviewed power shocks

Wentao Huang, Yixin Zhu and Weilin Yang*

School of Automation and Intelligent Science, Jiangnan University, Wuxi, China

* Corresponding author; E-mail: wlyang@jiangnan.edu.cn.

Highlights:

- A one-step barrier filter corrects the battery and supercapacitor power split.
- Each predicted margin must cover a fraction of the present reserve plus a fixed allowance.
- Of 500 shock trajectories, 488 met the sampled limits with the filter and 470 with matched projection.
- At a retention fraction of 0.80, the recovery intervals explain why all commands in the reserve-retention study came from fallback.

Abstract: A battery–supercapacitor system close to a storage limit can exhaust its reserve during a pulse missing from the demand forecast. We proposed a discrete-time barrier filter that corrects the power split from an energy-domain quadratic program at each 0.10-s update. It retains a fraction of each battery or supercapacitor margin and adds an allowance for uncertainty. The device power intervals give an exact feasibility check for the selected battery direction. Starting inside the nominal operating set, the model stays inside at subsequent samples if a feasible filter command is applied at every update. Under unpreviewed power shocks near an operating boundary, the filter kept 488 of 500 trajectories within the sampled physical limits (97.6%), compared with 470 (94.0%) under matched hard projection. These paired electrical simulations included losses, actuation delay and measurement error. Mean boundary dwell increased from 2.950 to 2.957 s. Guarded complementary filtering kept all 500 trajectories within limits and left less energy unserved. Its mean root-mean-square battery current was 65.10 A, compared with 28.28 A for the proposed controller. The nominal split accounts for most of that 56.6% reduction.

Keywords: hybrid energy storage system; battery–supercapacitor; model predictive control; discrete-time control barrier function; quadratic programming; safety filtering; uncertainty-aware control

1. Introduction

As electrification spreads across transport and industry, energy storage systems often have to supply a sustained load while covering brief peaks in power demand. A battery's high energy density makes it well suited to sustained loads, but its discharge current limit can leave a power shortfall during a peak on the direct-current (DC) bus. A battery–supercapacitor hybrid energy storage system (HESS) can draw the balance from the supercapacitor (SC), but only while the SC has energy to spare. Early HESS controllers

Copyright©2026 by the authors. Published by ELSP. This work is licensed under Creative Commons Attribution 4.0 International License, which permits unrestricted use, distribution, and reproduction in any medium provided the original work is properly cited.

assigned sustained demand to the battery and short pulses to the SC [1–4]. The available choices depend on the converter topology and its control [5,6]. If the SC starts near its lower voltage limit, power drawn early in a pulse leaves less for the remainder.

One option is to divert power to the battery as soon as the SC enters a guard band. Rule-based and complementary-filter schemes use this type of decision with little online computation [7,8]. In vehicle energy management, model predictive control (MPC) chooses the split over a finite horizon subject to device limits [9–11]. A soft state penalty still permits a violation. A hard limit can admit a command that uses the entire remaining margin in one sample. For the shock studied here, that command leaves the next bus request to a battery that may already be at its power limit.

A controller can track its current reference accurately even when the power request omits a shock. In Zhao *et al.* [12], a modified low-pass filter sets the split before multivariable feedback acts on current and bus voltage. At the vehicle level, Chen *et al.* used Koopman prediction for hierarchical vehicle following [13]. Ma and Li refined the demand forecast and tuned MPC weights and constraints with a genetic algorithm [14]. A pulse omitted even from the current preview entry is unavailable to the nominal allocator throughout the pulse.

Other HESS studies report results at different levels. Data-based microgrid control [15] and reinforcement-learning allocation [16] act on their own plant and demand models. Vehicle hybridization assessment [17] and converter validation [18] also use different objectives. We use same-plant comparisons below to measure the effect of reserve retention.

Here the operating region has four faces, two for battery charge and two for capacitor energy. Constrained MPC can enforce a limit on each predicted state [19], and robust MPC can move the limits inward to allow for bounded disturbances [20]. A control barrier function (CBF) describes the margin to a face [21,22]. A discrete-time control barrier function (DT-CBF) can require part of the current margin to remain at the next sample [23]. Previous work has imposed this condition inside MPC [24,25], combined several barriers [26], or treated high-order and adaptive barriers [27]. In HESS voltage and current regulation [28], the barrier acts below the power allocator considered here.

How much reserve can the supervisor retain once the shock reaches the bus? We start each of 500 paired trajectories 5% of the operating span from a boundary and leave the 8-s pulse out of the supervisor's preview. Over 200 samples per trajectory, we check battery state of charge (SOC) and both terminal voltages, integrate the power mismatch at the bus and count steps on emergency fallback. These measurements let us compare the gain in sampled safety with the cost of meeting the reserve requirement.

The nominal allocator is a convex quadratic program (QP) that uses battery SOC and SC energy as its states. A one-step filter corrects the power split. Its four affine inequalities give a battery power interval for the selected direction and a separate SC power interval. The limiting case is apparent in SOC units. At 50 kW the battery can change its nominal SOC by only about 4.1×10^{-5} per sample. A retention fraction of 0.80 would require an increase of about 3.90×10^{-3} at the reference lower-margin initialization. Reading the intervals in power units identifies requests the devices can actually supply.

Matched hard projection and uncertainty-aware DT-CBF act on the same nominal command in each shock realization. The nominal DT-CBF drops the uncertainty allowance. The offset-only filter drops the proportional retention term. These two cases separate the effects of the two terms on the admissible power

split. We also compare the guarded complementary filter (GCF), whose different allocation rule allows a separate comparison of battery current and unmet demand.

2. Technical background

2.1. HESS power allocation

In a fully active battery–SC topology, separate bidirectional converters track the battery and SC commands [2,3,6]. This permits independent supervisory power allocation, subject to the devices' SOC, voltage and power limits. Passive and semi-active topologies provide less independent control of the two branches.

A complementary filter sends the low-frequency power to the battery and the residual to the SC, with guard logic near a limit. In MPC the penalty on storage effort can change that split even for the same demand forecast [9,11]. For this reason we feed the same nominal power commands to the barrier and hard-projection filters when comparing their safety constraints.

2.2. Discrete-time barriers and constrained prediction

For a discrete system $x_{k+1} = f_d(x_k, u_k)$, a zeroing DT-CBF inequality has the generic form

$$h_i(x_{k+1}) \geq \rho_i h_i(x_k), \qquad 0 < \rho_i < 1, \tag{1}$$

where $h_i(x)$ is the signed margin to face i of the safe set. A positive margin is available reserve, and ρ_i is the fraction that must remain one sample later. The condition can be imposed at a prediction node in MPC or on the command issued by a nominal controller [24,25]. Starting with a nonnegative margin, applying a feasible command at every step keeps the model margin nonnegative by induction. Extending that result to the plant requires a bound on the error in the predicted next margin.

Our offset-only comparator uses the one-step constraint of Shaw Cortez *et al.* [29]. In their robust design the predicted margin covers an explicit disturbance bound. We use the design offsets from Section 3.3 and solve at every step. This comparison measures what proportional retention adds on the same plant. The original predictive scheme and its event trigger fall outside that comparison.

For the predictive comparison we use Didier and Zeilinger's approximate predictive control barrier function (APCBF) [30]. They approximate the predictive barrier with offline samples from the model to reduce online computation and analyze the effect of approximation and disturbance errors. Here both filters receive the same HESS model and the same nominal power request.

In energy coordinates, the nominal dynamics and the filter constraints are affine. An empty interval of admissible battery or SC power can therefore be traced to a physical power limit or to the chosen recovery requirement.

3. System and safety formulation

3.1. Two-model supervisory architecture

The evaluation plant includes resistance–capacitance (RC) polarization and equivalent series resistance (ESR), which are omitted from the affine controller model. Figure 1 shows how the controller and the electrical plant exchange commands and measurements. At sample k, the controller receives a four-step demand

preview and a shared state estimate. A nominal energy QP returns a preferred power split. The central comparison is between the matched operating-set hard-projection baseline and the uncertainty-aware DT-CBF projection. Supporting comparators include the unprojected nominal QP, physical hard projection, nominal DT-CBF, and GCF. The high-fidelity plant realizes the command through converter dynamics, electrical losses, and a delayed fast-residual split. All safety and service metrics are computed from the simulated electrical-plant trajectory.

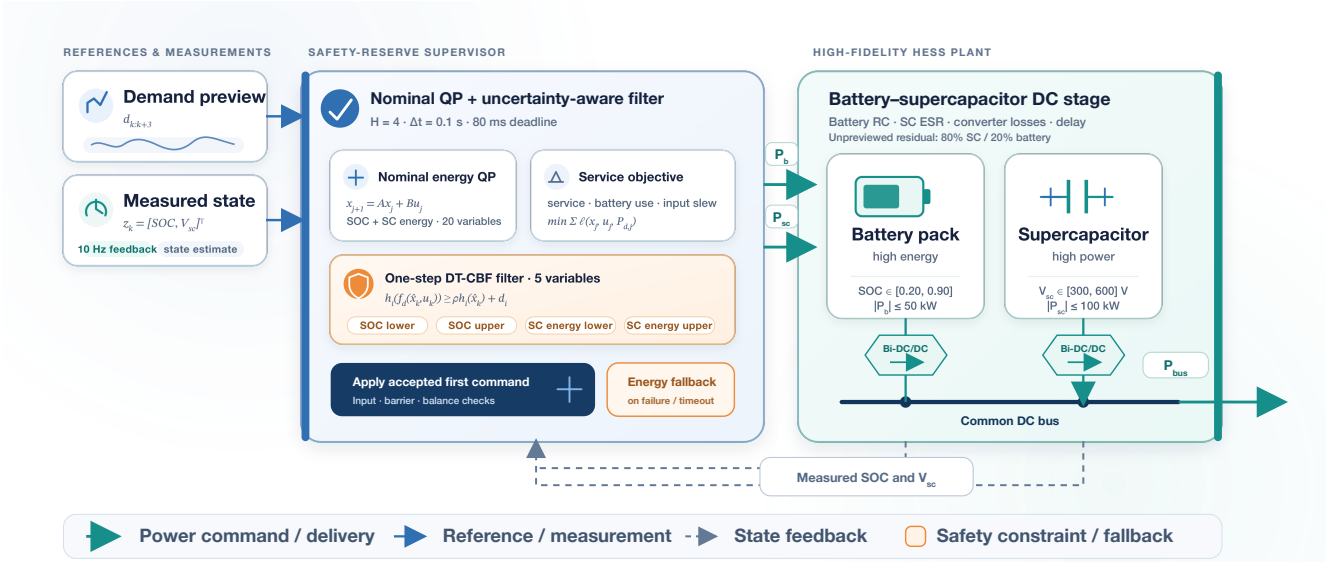

Figure 1. Supervisory controller and HESS plant. The four-step, 20-variable energy QP supplies the preferred allocation to the five-variable one-step filter. A command is accepted if it meets the input, barrier and power-balance conditions within the time limit. Otherwise, the common fallback supplies the command. The electrical plant includes RC/ESR dynamics, converter losses and delay, with an 80%/20% SC/battery split of the unpreviewed residual. Dashed links carry measured-state feedback. This figure was created in Gemini Image.

The controller state is

$$x_k = [\mathrm{SOC}_k,\ E_k]^\top, \qquad E_k = \tfrac{1}{2} C_{\mathrm{sc}} V_{\mathrm{sc},k}^2. \tag{2}$$

Positive storage power supplies the DC bus. Battery power is represented by nonnegative directional variables,

$$P_{\mathrm{b},k} = P_{\mathrm{b},k}^{\mathrm{dis}} - P_{\mathrm{b},k}^{\mathrm{chg}}, \tag{3}$$

which retain an affine SOC update over a sample Δt:

$$\mathrm{SOC}_{k+1} = \mathrm{SOC}_k - \frac{10^3 \Delta t}{E_{\mathrm{b}} \eta_{\mathrm{dis}}} P_{\mathrm{b},k}^{\mathrm{dis}} + \frac{10^3 \eta_{\mathrm{chg}} \Delta t}{E_{\mathrm{b}}} P_{\mathrm{b},k}^{\mathrm{chg}}, \tag{4}$$

$$E_{k+1} = E_k - 10^3 \Delta t P_{\mathrm{sc},k}. \tag{5}$$

Here E_{b} is the nominal battery energy in joules. The SC voltage is recovered from $V_{\mathrm{sc}} = \sqrt{2E/C_{\mathrm{sc}}}$ for reporting. The controller uses energy as its state.

3.2. Physical and operating sets

The controller represents physical energy limits by the box

$$\mathscr{S} = \left\{ x : \mathrm{SOC}_{\min}^{p} \le \mathrm{SOC} \le \mathrm{SOC}_{\max}^{p},\ E_{\min}^{p} \le E \le E_{\max}^{p} \right\}. \tag{6}$$

This energy box is distinct from the plant safety condition, which also requires battery and SC terminal voltages to remain in their physical intervals. In particular, capacitor voltage and terminal voltage differ under ESR current. The controller uses an interior reserve box $\mathscr{O} \subset \mathscr{S}$. The implemented SC guard combines an ESR estimate, a first-order delay allowance, and measurement error:

$$\Delta V_{\mathrm{op}} = \Delta V_{\mathrm{ESR}} + \Delta V_{\mathrm{delay}} + \bar{e}_V. \tag{7}$$

Thus $V^o_{\min} = V^p_{\min} + \Delta V_{\mathrm{op}}$ and $V^o_{\max} = V^p_{\max} - \Delta V_{\mathrm{op}}$, with the corresponding energy limits derived by Equation (2). The study uses the SOC interval $[0.240864, 0.859136]$ and capacitor-voltage reserve interval $[305.423, 594.577]$ V. These are design guards, not certified reachable-set bounds for the delayed plant. The operating-set hard-projection baseline varies a multiplier on the same guard in the safety–service–feasibility trade-off scenario (S3; Section 4.2).

3.3. Uncertainty offsets and the extent of the safety statement

Let $n_d = 2$, $c_Q = 0.10$, $c_R = 0.30$, $e_\eta = 0.03$, $\bar{e}_s = 0.01$, and $\bar{e}_V = 0.20$ V denote the declared design quantities. The filter uses the offsets

$$d_s = \bar{e}_s + \frac{10^3 P^{\max}_{\mathrm{b}} \Delta t}{E_{\mathrm{b}}}(c_Q + e_\eta + n_d), \tag{8}$$

$$d_E = C_{\mathrm{sc}} V^p_{\max} \bar{e}_V + 10^3 P^{\max}_{\mathrm{sc}} \Delta t (c_R + n_d), \qquad d = [d_s, d_s, d_E, d_E]^\top. \tag{9}$$

Thus $d_s = 0.0100822$ and $d_E = 35.0$ kJ. The first term in each expression converts measurement error to the corresponding state units. The remaining terms scale one step of nominal reserve consumption using the assumed capacity, efficiency, resistance and delay variations. The exact voltage-to-energy measurement-error bound is $CV^p_{\max}\bar{e}_V + \frac{1}{2}C\bar{e}_V^2$. Equation (9) keeps only its linear term. The resulting offsets set the required reserve. A certified plant bound would also have to cover RC/ESR dynamics, command history, the tails of the Gaussian preview noise and the unpreviewed pulse.

For each barrier face of the augmented plant, let e_i bound the margin error, $|h_i(\hat{x}_k) - h_i(x_k)| \le e_i$. Assume also that the realized next margin satisfies $h_i(x_{k+1}) \ge h_i(f_d(\hat{x}_k, u_k)) - w_i$. Both bounds must hold for every admissible history. Applying a command accepted by the filter then gives

$$h_i(x_{k+1}) \ge \rho h_i(x_k) + d_i - w_i - \rho e_i. \tag{10}$$

The sufficient inequality $d_i \ge w_i + \rho e_i$ would establish sampled energy-set preservation, conditional on feasibility and accepted commands at every step. Terminal-voltage preservation additionally needs a verified current/ESR map. For the electrical plant, neither these bounds nor recursive feasibility have been established. Its safety is assessed at the sampling instants in simulation.

3.4. Nominal QP and minimum-intervention filter

For a preview $d_{k:k+H-1}$, the nominal allocator minimizes

$$\begin{aligned} J_{\mathrm{nom}} = \sum_{j=0}^{H-1} \Big[w_\ell (P^2_{\mathrm{shed},j} + P^2_{\mathrm{spill},j}) + w_{\mathrm{b}} \big((P^{\mathrm{dis}}_{\mathrm{b},j})^2 + (P^{\mathrm{chg}}_{\mathrm{b},j})^2 \big) + w_{\mathrm{sc}} P^2_{\mathrm{sc},j} \Big] \\ + w_c \sum_{j=1}^{H} \big(\mathrm{SOC}_j - \mathrm{SOC}_c \big)^2 + w_\Delta \sum_{j=0}^{H-1} \| u_j - u_{j-1} \|_2^2, \end{aligned} \tag{11}$$

subject to Equations (4)–(5), directional battery bounds, SC power bounds, and the per-step bus balance

$$P_{\mathrm{b},j}^{\mathrm{dis}} - P_{\mathrm{b},j}^{\mathrm{chg}} + P_{\mathrm{sc},j} + P_{\mathrm{shed},j} - P_{\mathrm{spill},j} = d_j. \tag{12}$$

The nominal QP penalizes deviations of predicted SOC from its reference. SC energy enters the one-step filter but has no centering or terminal cost in the nominal QP. The increment term uses the applied prior command u_{-1} at the first prediction step. The sign of the previewed demand sets the battery's charge or discharge direction. A residual check rejects commands with non-negligible simultaneous charge and discharge.

The uncertainty-aware DT-CBF layer receives the first nominal command u_k^{nom} and solves

$$\begin{aligned} \min_{u_k, P_{\mathrm{shed},k}, P_{\mathrm{spill},k}} \quad & \|u_k - u_k^{\mathrm{nom}}\|_W^2 + w_\ell (P_{\mathrm{shed},k}^2 + P_{\mathrm{spill},k}^2) \\ \text{s.t.} \quad & (12), \quad u_k \in \mathscr{U}, \\ & h_i(x_k) + a_i^\top z_k - \rho h_i(x_k) \geq d_i, \quad i = 1, \ldots, 4, \end{aligned} \tag{13}$$

where $z_k = [P_{\mathrm{b}}^{\mathrm{dis}}, P_{\mathrm{b}}^{\mathrm{chg}}, P_{\mathrm{sc}}, P_{\mathrm{shed}}, P_{\mathrm{spill}}]^\top$, a_i follows from Equations (4)–(5), and d_i is the design allowance in Equations (8)–(9). The four barrier functions describe the lower and upper faces of $\mathscr{O}$:

$$\begin{aligned} h_1 &= \mathrm{SOC} - \mathrm{SOC}_{\min}^o, & h_2 &= \mathrm{SOC}_{\max}^o - \mathrm{SOC}, \\ h_3 &= E - E_{\min}^o, & h_4 &= E_{\max}^o - E. \end{aligned} \tag{14}$$

The nominal DT-CBF comparator sets $d_i = 0$ in Equation (13), whereas the proposed filter uses the uncertainty offsets. The physical hard-projection comparator constrains the physical next state, while the matched hard-projection baseline constrains the next state in the operating set.

Proposition 1: Let $x_k \in \mathscr{O}$ be the exact state supplied to the filter under the nominal dynamics in Equations (4)–(5). Applying an accepted uncertainty-aware DT-CBF command that satisfies Equation (13) gives $h_i(x_{k+1}) \geq \rho h_i(x_k) + d_i \geq 0$ on every face. If such a command is applied at each subsequent step, the state remains in $\mathscr{O}$ at every sampling instant.

Proof: The energy-state dynamics are affine, so the barrier-constraint residual in Equation (13) is exactly $h_i(x_{k+1}) - \rho h_i(x_k) - d_i$. Since $h_i(x_k) \geq 0$, $\rho \in (0,1)$, and $d_i \geq 0$, the accepted command gives $h_i(x_{k+1}) \geq 0$. Induction over samples establishes the stated result.

Proposition 1 applies to energy margins at the sampling instants of the nominal model. In the electrical plant, RC/ESR dynamics, conversion losses, delayed actuation and measurement error can alter the relationship between predicted energy margins and realized terminal-voltage margins. The offsets are design allowances rather than bounds on this combined mismatch, so safety under these effects is evaluated through the simulations.

3.5. Feasible recovery intervals and relaxation

For the SC, the four-face filter reduces to the interval

$$\max\left\{-P_{\mathrm{sc}}^{\max}, \frac{d_E - (1-\rho)h_4}{10^3 \Delta t}\right\} \leq P_{\mathrm{sc}} \leq \min\left\{P_{\mathrm{sc}}^{\max}, \frac{(1-\rho)h_3 - d_E}{10^3 \Delta t}\right\}. \tag{15}$$

For a fixed battery direction, let κ be the corresponding positive SOC decrement per kW from Equation (4), and let $[P_b^-, P_b^+]$ be $[0, 50]$ for discharge or $[-50, 0]$ for charge. The battery interval is

$$\max\left\{P_b^-, \frac{d_s-(1-\rho)h_2}{\kappa}\right\} \le P_{\mathrm{b}} \le \min\left\{P_b^+, \frac{(1-\rho)h_1-d_s}{\kappa}\right\}. \tag{16}$$

With unbounded nonnegative shedding/spillage variables, bus balance does not repair an empty battery or SC interval. These two interval intersections therefore characterize model feasibility exactly for the selected directional branch.

Near the lower SOC face during a positive-demand branch, the battery cannot charge. Even zero battery power then requires $(1-\rho)h_1 \ge d_s$. At the nominal initialization of the unpreviewed power-shock scenario (S2; Section 4.2), $h_1 = 0.05(0.859136-0.240864) = 0.030914$, giving $\rho \le 1-d_s/h_1 \simeq 0.674$ before measurement error. The symmetric restriction applies at the upper face on a charging branch. This identifies a structural source of failure at high ρ: the rate condition asks for inward recovery that the fixed battery direction forbids. Even when the direction permits recovery, the largest nominal SOC change at 50 kW is only about 4.1×10^{-5} per sample. At the reference initial margin and $\rho = 0.80$, the required additional margin is about 3.90×10^{-3}, far above that actuator capability. The SC power interval and subsequent state-estimation errors can restrict feasibility further. An SOC offset based on measurement error can thus exceed what the battery can recover in one sample, even well inside the physical SOC limits.

A slackened inequality $h_i(f_d) \ge \rho h_i(\hat{x}) + d_i - \xi_i$, $\xi_i \ge 0$, can remove this algebraic infeasibility when penalized slacks are unrestricted. The original preservation implication then loses ξ_i in its lower bound. If only the rate term is relaxed, the hard constraint $h_i(f_d) \ge 0$ still enforces the nominal box whenever that problem is feasible. It does not bound terminal voltages under delayed actuation. The offset-only comparator below tests a less restrictive alternative by setting the retention term to zero without introducing a tuned slack penalty.

3.6. Common energy-bounded fallback

All methods use the same energy-bounded fallback. If optimization fails or exceeds its deadline, a 50/50 battery–SC allocation is projected onto one-step directional power intervals determined by the current operating-set energy margins. The remaining imbalance is assigned to P_{shed} or P_{spill}. For example, the available SC discharge power is

$$\bar{P}_{\mathrm{sc}}^{\mathrm{dis}} = \min\left\{P_{\mathrm{sc}}^{\max}, \frac{E_k - E_{\min}^{o}}{10^3\Delta t}\right\}. \tag{17}$$

The other endpoints follow Equations (4)–(5). Fallback is triggered by nominal-QP failure, projection primal infeasibility, an iteration limit, non-negligible simultaneous charging and discharging, or computation beyond 80 ms, including converged but late solutions. These commands respect nominal directional energy bounds but carry no separate terminal-safety guarantee for the delayed plant.

4. Simulations

4.1. Control-oriented parameterization and plant

Table 1 gives the synthetic parameters used for the controller and the electrical plant. The battery model has one Thevenin RC polarization state and separate charge and discharge efficiencies. The SC model includes series resistance and leakage. Each converter tracks a saturated command with first-order

dynamics, 97% efficiency and a delay of up to two samples, including zero delay. We add bounded errors to the SOC and voltage observations. After each update, the simulated SOC and voltage retain their computed values, including any excursions beyond the physical limits.

Table 1. Control-oriented HESS and simulation parameters.

Quantity	Symbol	Value
Sampling period/deadline	Δt, T_d	0.10 s/0.08 s
Nominal battery capacity	Q_{b}	100 Ah
Battery open-circuit voltage	V_{oc}	300–400 V
Battery power limit	$\lvert P_{\mathrm{b}}\rvert$	50 kW
Battery physical SOC range	SOC	0.20–0.90
SC capacitance	C_{sc}	100 F
SC physical voltage	V_{sc}	300–600 V
SC power limit	$\lvert P_{\mathrm{sc}}\rvert$	100 kW
Converter efficiency	η_{conv}	0.97
ESR/leakage	R_{sc}, γ_{leak}	0.01 Ω, 10^{-5} s^{-1}
Nominal QP horizon	H	4 samples
Uncertainty-aware DT-CBF rate	ρ	0.50 (unless scanned)

4.2. Uncertainty and scenarios

The simulations use a common uncertainty model. At normalized mismatch level λ, plant resistance can rise by up to 0.30λ, usable capacity can decrease by up to 0.10λ, and each battery efficiency can decline by up to 0.03λ. The converter time constant receives an additional 0.04λ s, the discrete delay is sampled from zero to two steps, preview bias is bounded by $\pm 8\lambda$ kW with 1.5λ kW noise standard deviation, and measurement errors are bounded by 0.01λ in SOC and 0.20λ V in SC voltage. The robust operating set and barrier offsets are based on these specified parameter bounds.

The four simulation scenarios examine normal mismatch (S1), unpreviewed power shocks (S2), safety–service–feasibility trade-offs (S3), and serial computational performance (S4).

S1 contains five mismatch levels $\lambda \in \{0, 0.25, 0.50, 0.75, 1.0\}$ and 100 paired realizations per level. It begins sufficiently far from the operating boundary to establish the normal-mismatch reference. S2 fixes $\lambda = 1$ and comprises 500 paired 20-s trajectories. Each S2 initial condition lies 5% of the operating-set span from either the lower or upper boundary. A sign-matched, unpreviewed 70-kW DC-bus pulse lasting 8 s is inserted at a uniformly drawn time between 6 and 12 s. Throughout the pulse, the lower bus loop sees the realized demand and uses the SC to cover the residual. The supervisor sees the measured storage states. Its current preview entry and all later entries omit the pulse for the full 8 s, so the change in reserve is the supervisor's feedback on the shock.

S3 compares five operating-set hard-projection margin multipliers (0.50, 0.75, 1.00, 1.25, and 1.50) and five uncertainty-aware DT-CBF rates (0.50, 0.80, 0.90, 0.95, and 0.99) using the same 500 paired S2 realizations. S4 measures whole-controller elapsed time at the default settings under the same S2 conditions, using serial execution.

The parameter study also evaluates $\rho = 0.60$ and 0.70 and an offset-only safety filter on the same 500 paired S2 realizations. Eight held-out 20-s shock families (seeds 40,000–40,049) examine different pulse magnitudes, durations, directions and initial reserves: 40 kW/4 s, 100 kW/8 s, 70 kW/2 s, 70 kW/12 s, opposite-sign 70 kW/8 s, two same-sign 70-kW/3-s pulses, two alternating-sign 70-kW/3-s pulses, and

mixed initial reserves with a 70-kW/8-s pulse. Single pulses begin at 6 s. For repeated shocks, the two pulses occur at 6–9 and 12–15 s. In the mixed-reserve family, the initial SOC cycles through fractions 0.15/0.50/0.85 of its operating interval and the initial SC energy through fractions 0.02/0.20/0.50/0.80/0.98 of its interval. The pulse sign alternates independently. Other families retain S2's paired upper/lower initialization and matched charge/discharge sign, except the explicitly opposite-sign family. Realized demand changes while preview remains fixed. All five principal policies were evaluated on the same scenarios without retuning the weights, offsets or GCF coefficient.

4.3. Comparison methods and performance metrics

Table 2 lists the controllers. All methods used the same demand realizations, previews, initial plant states and observation model. They also shared the physical input limits, 80-ms budget and analytical fallback. Corresponding QP weights are the same wherever the objectives are shared. Matched hard projection requires the predicted state to lie in the operating set. Uncertainty-aware DT-CBF additionally constrains the fraction of each margin retained and includes the uncertainty offset.

GCF updates its battery reference as $\bar{P}_{b,k} = 0.8\bar{P}_{b,k-1} + 0.2\hat{d}_k$, starting at zero, clips it to the 50-kW input interval, and assigns the preview residual to the SC. This nominal split then passes through the same operating-set hard-projection QP as the matched baseline, with five decision variables. Like the other supervisors, GCF uses the current preview entry $\hat{d}_k$ without an additional measurement of actual demand.

Table 2. Compared allocation and safety policies.

Policy
Unprojected nominal energy QP.
One-step physical-set hard projection.
Matched operating-set hard projection.
Nominal DT-CBF projection, $d_i = 0$.
Uncertainty-aware DT-CBF projection, with the design offsets in Equations (8)–(9).
Complementary filter with matched operating guards.
Offset-only one-step filter, $h_i(f_d(\hat{x},u)) \geq d_i$, adapted from the $N = 1$ condition in [29]; the same nominal QP, offsets, directional branch, objective and fallback are used, without its event trigger or a certified plant disturbance bound.

We check plant SOC, battery terminal voltage and SC terminal voltage after each of the 200 state updates. A trajectory passes if all three stay within their physical limits at every check. Excursions between samples lie outside this criterion. For each channel y, we record the depth $\max\{y_{\min} - y, y - y_{\max}, 0\}$ in its native units. Its violation duration equals the number of out-of-limit samples times Δt.

Boundary dwell sums the time SOC or SC terminal voltage lies inside its physical interval but outside its operating interval:

$$T_{\text{dwell}} = \Delta t \sum_k \sum_{y \in \{\text{SOC}, V_{\text{sc,term}}\}} \mathbf{1}\{y_k \in [y^p_{\min}, y^p_{\max}] \setminus [y^o_{\min}, y^o_{\max}]\}. \tag{18}$$

If both variables lie in their respective boundary bands at the same sample, each contributes to the dwell time. A variable outside its physical limits contributes no dwell time at that sample. All policies and S3 guard multipliers are evaluated against the same default operating box. Scheduled shedding energy is $\Delta t \sum_k P_{\text{shed,k}}/3600$ in kWh, using kW power. It records requested shedding, not an additional physical

energy loss. With actual bus mismatch $e_k = P_{\mathrm{d},k} - P_{\mathrm{b},k}^{\mathrm{real}} - P_{\mathrm{sc},k}^{\mathrm{real}}$, total positive unserved energy is

$$E_{\mathrm{unserved}} = \frac{\Delta t}{3600} \sum_k \max\{e_k, 0\}. \tag{19}$$

Rejected regenerative energy and excess injection both contribute to the negative part of the mismatch, $E_- = \Delta t \sum_k \max\{-e_k, 0\}/3600$. We add it to the positive unserved energy to obtain $E_{\mathrm{abs}} = E_{\mathrm{unserved}} + E_-$. This accounts for regenerative service as well as unmet load. The measured bus powers already reflect delayed shedding or spillage, so adding the requested amounts would count them twice. For each trajectory we count the N_f control steps that used the fallback and report $100N_f/N$ with $N = 200$. Solver status and elapsed time identify the trigger.

Battery current is reconstructed from the plant's exact SOC increments and its sampled capacity and directional efficiency, including the initial SOC:

$$I_{b,k} = \begin{cases} -\Delta \mathrm{SOC}_k Q \eta_{\mathrm{dis}}/\Delta t, & \Delta \mathrm{SOC}_k \le 0, \\ -\Delta \mathrm{SOC}_k Q/(\eta_{\mathrm{chg}} \Delta t), & \Delta \mathrm{SOC}_k > 0. \end{cases} \tag{20}$$

Here Q is in ampere-seconds. We report per-trajectory root-mean-square (RMS) and peak current, SOC range including the initial state, RMS of $(P_{b,k}^{\mathrm{real}} - P_{b,k-1}^{\mathrm{real}})/\Delta t$ (initial power zero), and $\Delta t \sum I_{b,k}^2 R_0$. The last quantity is ohmic energy in joules, not temperature or a calibrated aging model. Each tabulated metric is first calculated for each trajectory and then averaged. In particular, the peak-current column gives the mean of the individual trajectory peaks.

5. Results

5.1. Reference behavior away from the operating boundary

S1 provides the reference condition in which normal load variation and mismatch leave most of the operating reserve available. The matched hard projection and uncertainty-aware DT-CBF both completed all the trajectories safely at each mismatch level, with identical boundary-dwell and scheduled-shedding outcomes. Neither safety layer materially changed the allocation while the operating constraints were inactive.

5.2. Unpreviewed power-shock response

The hard policy constrains the nominal next state in a fixed operating box. The DT-CBF condition additionally retains a proportion of the current margin and adds its design allowance, restricting reserve depletion before the nominal next state reaches that boundary.

Figure 2 illustrates the distinction between fixed-boundary and rate-shaped margin policies. Its curves are schematic: they do not imply a guaranteed ordering of physical trajectories or a measured dwell reduction.

Figure 3 shows the response to the same S2 shock and observation noise under both policies (seed 20,008). In the inset in panel (b), the minimum SC terminal-voltage margin was $+0.050$ V with uncertainty-aware DT-CBF and -0.090 V with matched hard projection. Only the latter crossed the physical limit at a sampling instant. Panel (c) shows the corresponding SC commands. Panel (d) gives the distribution of boundary dwell time over all 500 paired trajectories.

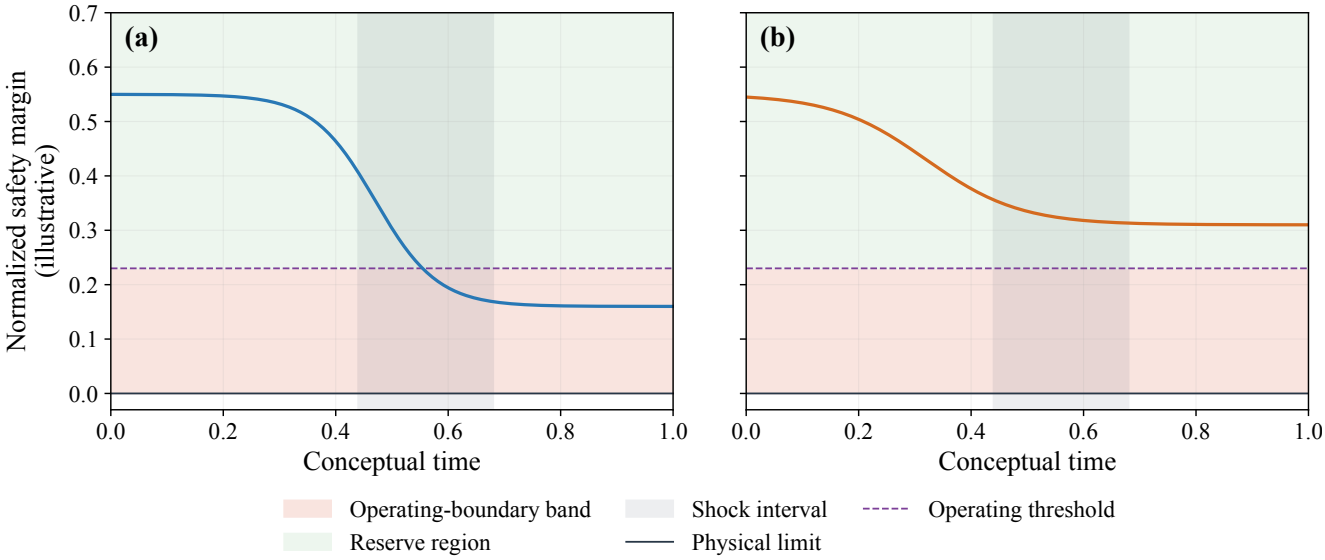

Figure 2. Illustrative margin geometry for fixed operating-boundary and rate-shaped DT-CBF policies. The curves depict possible margin evolution, not simulation trajectories or a performance comparison. **(a)** Fixed operating-boundary constraint; **(b)** Rate-shaped DT-CBF condition.

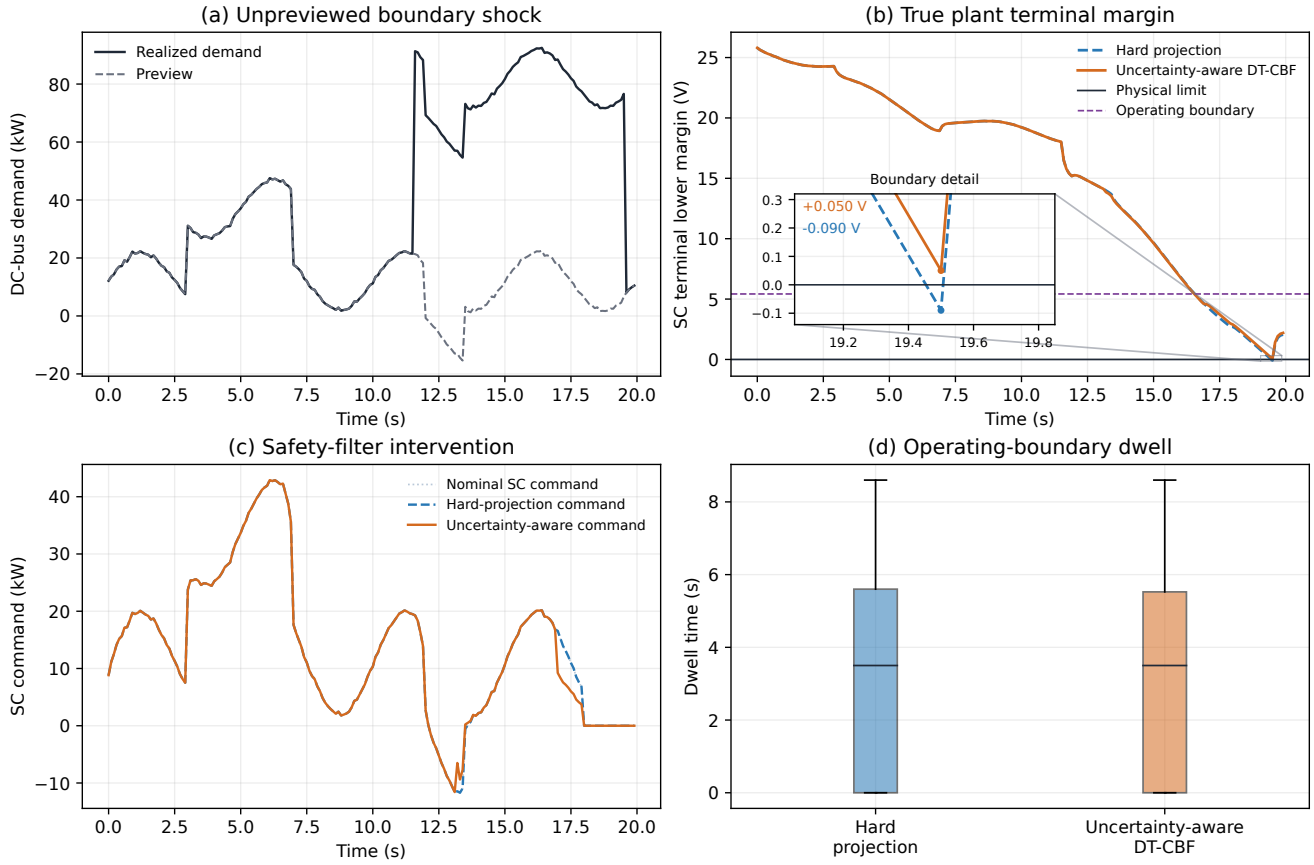

Figure 3. S2 power-shock response and boundary-dwell comparison. Panels **(a)–(c)** show a local safety-benefit case (seed 20008). **(a)** Realized demand contains an unpreviewed pulse relative to the controller preview; **(b)** True SC terminal-margin trajectories for the same shock realization, with a boundary-detail inset; **(c)** Nominal and filtered SC power commands; **(d)** Operating-boundary dwell over all 500 paired S2 runs, not only the illustrated case.

At the default settings, 488/500 trajectories stayed within the sampled physical limits with uncertainty-aware DT-CBF, compared with 470/500 with matched hard projection (Table 3). The corresponding mean dwell times were 2.957 and 2.950 s. Thus, more trajectories met the safety criterion without a reduction in mean dwell time. Both comparisons include steps supplied by fallback.

Table 3. Safety and service under 500 paired S2 power shocks. Scheduled shedding requests and total physical unserved energy are in kWh. "Sched." below means scheduled shedding requests.

Method	Safe/500	Dwell(s)	Sched.	Total	Fallback (%)
Nominal QP	353	2.895	0.000001	0.002197	0.000
Physical projection	353	2.895	0.000000	0.002197	0.004
Matched projection	470	2.950	0.003222	0.005023	11.421
Nominal DT-CBF	471	2.950	0.003222	0.005023	11.464
Uncertainty-aware DT-CBF	488	2.957	0.003160	0.004963	14.379
GCF	500	0.000	0.000000	0.002303	0.000
Offset-only filter	479	2.953	0.003196	0.004998	12.862

5.3. All-controller comparison

Of the controllers in Table 3, the unprojected nominal QP and physical hard projection each kept 353 trajectories within the sampled limits. The corresponding counts were 471 for nominal DT-CBF and 488 for uncertainty-aware DT-CBF. GCF met the limits on every trajectory, with zero boundary dwell and no scheduled shedding. Its delivered power still differed from the realized demand, leaving some energy unserved. Within the predictive controllers, the gain in sampled safety with the uncertainty-aware filter came with more frequent fallback than under matched projection or nominal DT-CBF.

Table 4 shows the difference in battery loading. Mean RMS battery current was 65.10 A with GCF and 28.28 A with uncertainty-aware DT-CBF. The corresponding mean ohmic energies were 3034 and 579 J. The DT-CBF controller also gave a lower mean peak current, a smaller SOC swing and a lower RMS rate of change of battery power (12.62 *versus* 16.03 kW/s). Mean time at $|P_b| \geq 49.5$ kW was 0.064 s for GCF and 0.000 s for the proposed policy. These short mean durations give no indication of prolonged battery saturation. Predictive allocation reduced battery loading, whereas GCF left less demand unserved and kept more trajectories within the limits. Neither temperature rise nor cycle life was evaluated.

Table 4. S2 battery metrics over 500 paired trajectories. Entries are mean per-trajectory values; Peak is the average of the peak current in each trajectory.

Method	RMS (A)	Peak (A)	SOC swing (%)	Slew (kW/s)	I^2R (J)
Nominal QP	26.42	49.48	0.110	12.30	495
Physical projection	26.42	49.51	0.110	12.31	495
Matched projection	27.63	57.22	0.119	12.57	548
Nominal DT-CBF	27.64	57.27	0.119	12.57	549
Uncertainty-aware DT-CBF	28.28	59.59	0.121	12.62	579
GCF	65.10	137.14	0.313	16.03	3034
Offset-only filter	27.97	58.48	0.120	12.59	564

Nominal DT-CBF completes 471/500 trajectories safely, compared with 488/500 for the offset counterpart. Fallback rates are 11.46% and 14.38%, while total unserved energies are 0.005023 and

0.004963 kWh. In S2, adding the offset increased sampled safety and fallback use, with little change in total unserved energy.

The 147 unsafe nominal-QP trajectories and the 147 unsafe physical-projection trajectories are SC terminal-voltage violations, with worst depths of 0.963 and 0.962 V, respectively. Neither set contains SOC or battery-terminal violations. Their QPs remain feasible with essentially zero scheduled shedding because they constrain nominal allocations, whereas the unseen residual, converter dynamics and ESR determine plant terminal voltage. A feasible optimization explains why essentially no fallback is required despite physical violations. The small physical-projection fallback count is due to late computation, not detection of those terminal-voltage violations. Other methods' native-unit severity, duration and trigger counts are reported in Table 5.

Table 5. S2 violation severity and dominant fallback trigger; counts span 100,000 steps per method.

Method	SC unsafe	Worst (V)	Duration (s)	Triggers (I/K/R/L)
Nominal QP	147	0.963	0.0626	0/0/0/0
Physical projection	147	0.962	0.0624	0/0/0/4
Matched projection	30	0.361	0.0080	11404/16/1/0
Nominal DT-CBF	29	0.361	0.0078	11424/33/6/1
Uncertainty-aware DT-CBF	12	0.224	0.0026	14347/24/5/3
GCF	0	0.000	0.0000	0/0/0/0
Offset-only filter	21	0.314	0.0048	12836/19/2/5

All SOC and battery-terminal violation counts are zero. I: primal infeasibility; K: iteration limit; R: constraint-residual rejection; L: computational deadline exceeded, including converged solutions. Violation durations are averaged over all 500 trajectories.

5.4. Reserve retention and recovery feasibility

Figure 4 compares the S3 design points. With the widest hard-projection margin, all trajectories met the sampled safety criterion. The fallback rate was 15.05%, mean dwell time was 2.940 s and mean unserved energy was 0.006171 kWh. Uncertainty-aware DT-CBF at $\rho = 0.50$ kept 488/500 trajectories within the limits, with 14.38% fallback, 2.957 s mean dwell and 0.004963 kWh mean unserved energy. The wider fixed margin thus kept more trajectories safe but left more energy unserved than the filter at this setting.

At high retention rates ($\rho \geq 0.80$), fallback supplies all commands (Table 6). Safety and service at these points therefore characterize the fallback rather than the DT-CBF filter.

Table 7 reports the corresponding hard-projection margin sweep, using the same boundary-dwell reporting band at every setting.

The intermediate rates in Table 6 show that the transition is already substantial below 0.80: fallback is 14.38% at $\rho = 0.50$, 33.20% at 0.60 and 69.16% at 0.70. Their different dwell and unserved-energy values include the behavior of the common fallback. Consequently, a favorable high-rate energy value cannot be attributed to accepted barrier commands alone. Equations (15) and (16) explain the high-rate failure through empty recovery intervals rather than a nonconvex local optimizer. A feasible retention rate must leave both power intervals nonempty for the battery's chosen direction.

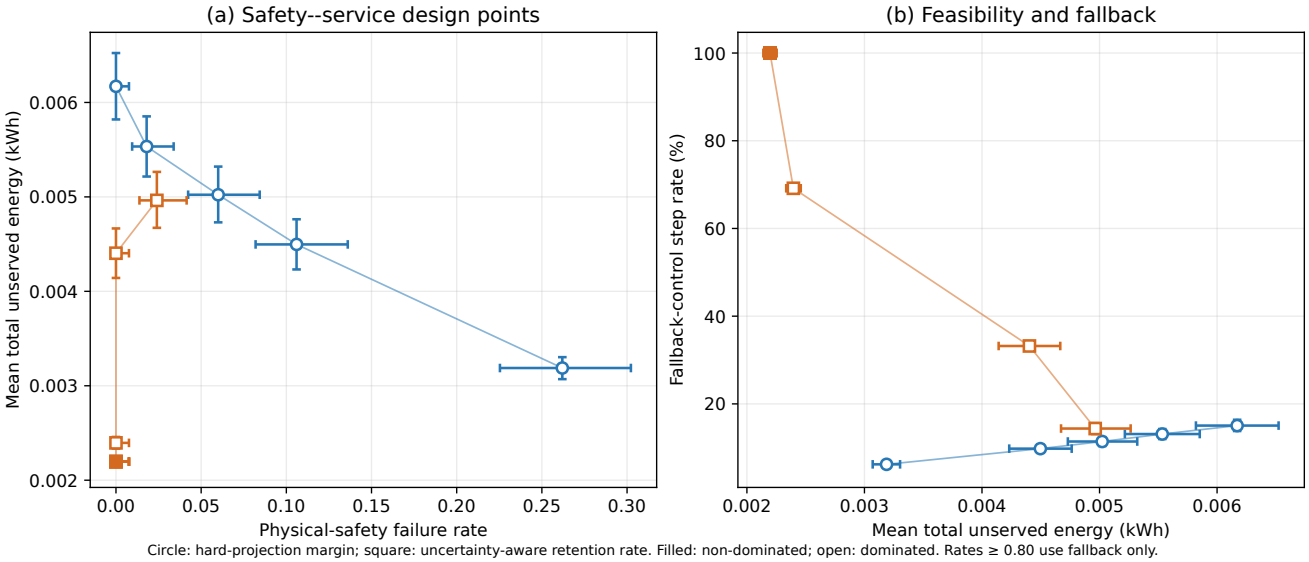

Figure 4. S3 reserve-policy comparison. **(a)** Physical-safety failure rate *versus* mean total unserved energy over 500 shock realizations. Circles denote hard-projection margin multipliers and squares denote uncertainty-aware barrier rates. Filled markers are non-dominated among the declared design points; **(b)** Total unserved energy *versus* fallback rate. Points at $\rho \geq 0.80$ use fallback exclusively. Error bars in **(a)** indicate uncertainty in the estimated failure rate and mean energy (95% Wilson and bootstrap intervals, respectively; 10,000 resamples for energy).

Table 6. Effect of reserve retention on safety, service and fallback over 500 paired S2 scenarios.

ρ	Safe/500	Dwell (s)	Total (kWh)	RMS (A)	Fallback (%)
0.50	488	2.957	0.004963	28.28	14.38
0.60	500	2.600	0.004404	31.99	33.20
0.70	500	1.205	0.002395	37.25	69.16
0.80	500	0.331	0.002197	40.53	100.00
0.90	500	0.331	0.002197	40.53	100.00
0.95	500	0.331	0.002197	40.53	100.00
0.99	500	0.331	0.002197	40.53	100.00

Table 7. S3 hard-projection margin comparison over 500 paired realizations, with a common reference band for boundary dwell.

Margin	Safe/500	Dwell (s)	Total (kWh)	Fallback (%)
0.50	369	2.901	0.003188	6.22
0.75	447	2.941	0.004497	9.79
1.00	470	2.950	0.005023	11.42
1.25	491	2.956	0.005534	13.13
1.50	500	2.940	0.006171	15.05

5.5. Offset-only filtering and held-out shocks

The offset-only filter kept 479/500 S2 trajectories within the sampled limits, with 2.953 s mean dwell, 12.86% fallback and 0.004998 kWh total unserved energy. Adding proportional retention at $\rho = 0.50$ raised the count to 488/500 (Table 3). The offset-only comparison uses the $N = 1$ constraint of Shaw Cortez *et al.* [29], but not their certified plant disturbance bound, event trigger or full predictive controller.

Table 8 summarizes sampled safety across the additional shock families. Tables 9 and 10 report positive unserved energy and negative bus-mismatch energy, respectively. For both the 100-kW/8-s and the 70-kW/12-s shocks, each DT-CBF filter kept 29/50 trajectories within the physical limits at the sampling instants. GCF did so on every trajectory in both families. With mixed initial states, the count was 44/50 for all three policies.

In S2, 17 trajectories met the sampled safety criterion with the uncertainty-aware filter but not with nominal DT-CBF. None showed the reverse outcome. In the 400 additional paired scenarios, the two filters agreed on every trajectory. GCF alone met the criterion in 42 of those cases. The fixed offsets therefore prevented some sampled violations in S2 but gave no additional safe trajectories in the other shock families.

Table 8. Sampled physical-safety success under eight held-out shock families (out of 50).

Family	HP	N	U	GCF	O
mild	50	50	50	50	50
severe	29	29	29	50	29
short	50	50	50	50	50
long	29	29	29	50	29
opposite	50	50	50	50	50
repeated	50	50	50	50	50
alternating	50	50	50	50	50
mixed initial	44	44	44	44	44

Each cell uses 50 paired held-out seeds. HP: matched projection; N: nominal DT-CBF; U: uncertainty-aware DT-CBF; O: offset-only filter.

Table 9. Mean total unserved energy under held-out shocks (Wh per trajectory).

Family	HP	N	U	GCF	O
mild	1.85	1.85	1.85	2.12	1.85
severe	9.03	9.03	8.99	7.97	9.01
short	2.17	2.17	2.17	3.11	2.17
long	5.83	5.83	5.83	4.06	5.83
opposite	2.18	2.18	2.18	3.47	2.18
repeated	2.83	2.83	2.83	3.77	2.83
alternating	2.78	2.78	2.78	3.72	2.78
mixed initial	4.13	4.13	4.09	4.86	4.11

Each cell uses 50 paired held-out seeds. HP: matched projection; N: nominal DT-CBF; U: uncertainty-aware DT-CBF; O: offset-only filter. Regenerative rejection is a separate service channel.

Table 10. Mean negative physical bus-mismatch energy under held-out shocks (Wh), including rejected regenerative absorption.

Family	HP	N	U	GCF	O
mild	1.74	1.74	1.74	1.71	1.74
severe	2.08	2.08	2.08	2.15	2.08
short	2.06	2.06	2.06	1.95	2.06
long	1.87	1.87	1.87	1.78	1.87
opposite	2.07	2.07	2.07	1.91	2.07
repeated	2.72	2.72	2.72	2.61	2.72
alternating	2.67	2.67	2.67	2.56	2.67
mixed initial	1.99	1.99	2.00	2.00	2.00

Each cell uses 50 paired held-out seeds. HP: matched projection; N: nominal DT-CBF; U: uncertainty-aware DT-CBF; O: offset-only filter.

For the comparison with Didier and Zeilinger's APCBF [30], we gave both filters the same HESS, nominal allocator, intervention objective and 10-s computational budget. Both could command battery charging or discharging. The 80-ms deadline applies to the separate timing study.

Without shock feedforward, APCBF kept all 500 S2 trajectories within the sampled limits and uncertainty-aware DT-CBF kept 488. Mean absolute bus-mismatch energy was 4.276 Wh with APCBF and 6.759 Wh with DT-CBF. When both supervisors received the current bus demand but no future shock preview, both passed all 500 trajectories. DT-CBF then had the smaller mismatch (4.302 *versus* 4.537 Wh), although it used the common emergency fallback on 12.36% of steps and APCBF used it on none. In the additional shock cases, safety counts matched under both information conditions. APCBF had lower mean absolute mismatch energy in both.

The APCBF approximation is trained on state–barrier samples generated by the plant model. This avoids the need for historical operating records but still requires sampling and validating the region near the safety boundary. Changes to the plant parameters or constraints may require retraining. The affine filter sets retention directly through ρ. Both approaches require a plant model and a treatment of uncertainty.

5.6. Computational performance

Figure 5 compares the elapsed-time statistics, deadline-overrun rates and fallback rates of matched projection and uncertainty-aware DT-CBF. For the uncertainty-aware controller, the mean, 99th-percentile (p99), 99.9th-percentile (p99.9) and maximum elapsed times were 1.078, 1.233, 1.599 and 8.159 ms, respectively. The corresponding times for matched projection were 1.076, 1.228, 1.573 and 31.120 ms. Both controllers met the 80-ms deadline at all 100,000 steps on the test computer. Execution and scheduling times on embedded hardware remain to be measured.

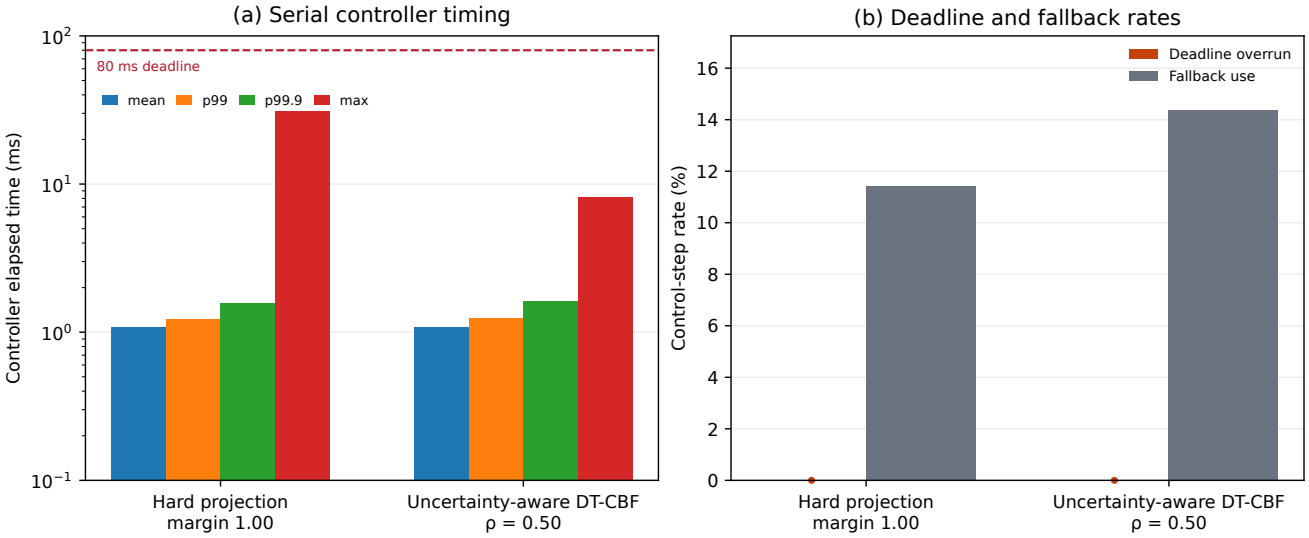

Figure 5. S4 computational performance over 500 S2 scenarios per policy. **(a)** Whole-controller elapsed time and the 80-ms computational deadline; p99 and p99.9 are the 99th and 99.9th percentiles; **(b)** Deadline-overrun and fallback rates per control step.

We used Python and the Operator Splitting Quadratic Program (OSQP) solver [31] on an Apple M1, with separate instances for the 20-variable nominal QP and the five-variable filter. Both used absolute and relative tolerances of 10^{-7}, a 20,000-iteration cap, scaling level 10 and no polishing. The alternating

direction method of multipliers (ADMM) started with a penalty of 0.1, adapted it during solution, warm-started the primal variables and checked termination every 25 iterations. We scaled the SOC and energy rows by their respective interval spans.

In S2, S3 and the additional shock comparisons, we checked the reconstructed signed command before applying it. A normalized barrier residual below -10^{-7} triggered fallback. So did an input or service bound violation above 10^{-5} kW, or an absolute power-balance error above the same threshold. The nominal preservation proof uses exact inequalities.

Timing covered the 500 S2 scenarios in serial execution. OpenBLAS is an optimized Basic Linear Algebra Subprograms (BLAS) library, and OpenMP denotes Open Multi-Processing. The OpenBLAS, OpenMP and vecLib runtimes were each limited to one thread. Elapsed time included matrix construction, solver initialization, both QPs, feasibility checks and fallback when required. State observation, plant propagation, plotting and disk input/output were excluded.

Figure 6 returns to the shock responses, showing paired trajectories for the local S2 safety benefit and a held-out counterexample.

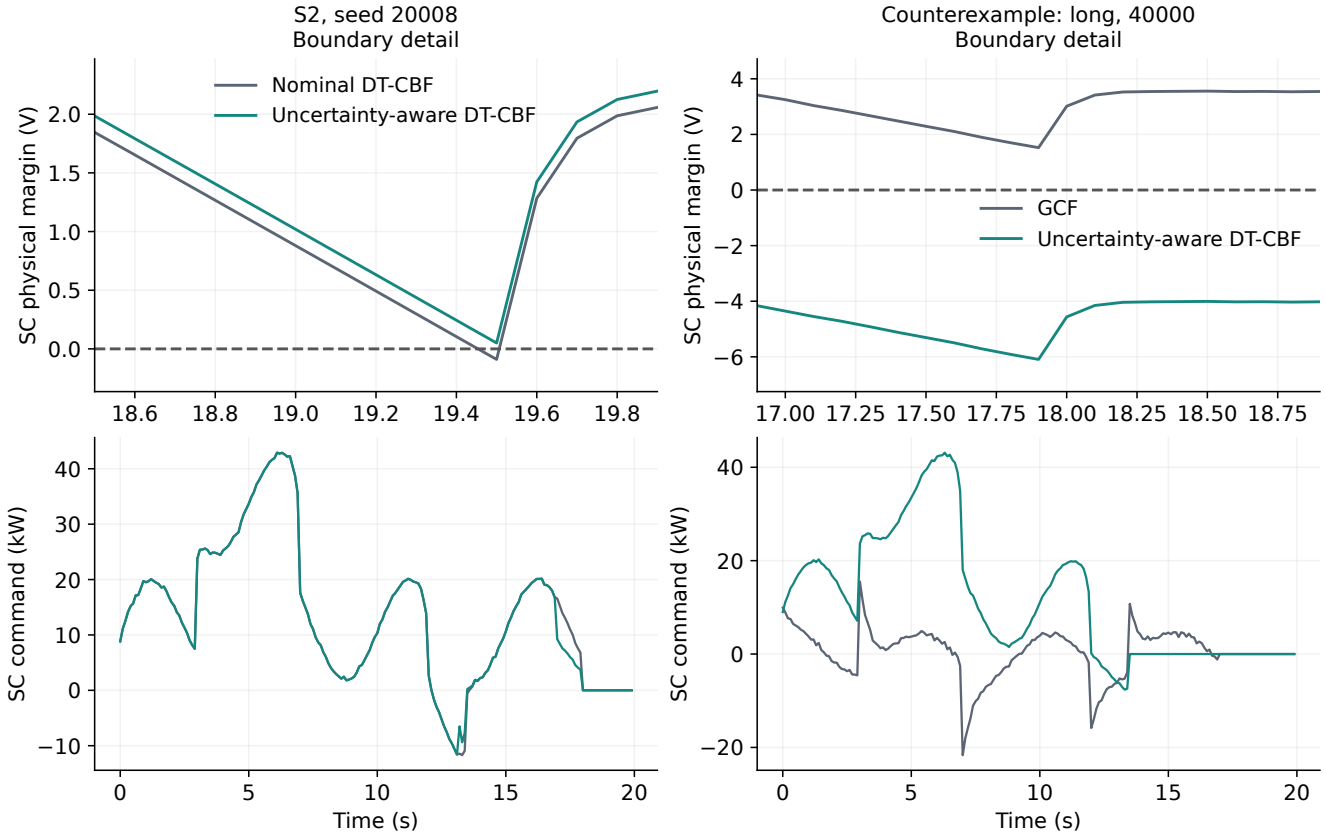

Figure 6. Matched shock responses illustrating local safety improvement and a limit of the fixed reserve policy. The first S2 seed where the uncertainty-aware policy is safe and nominal DT-CBF is unsafe is used in the first column. The second column uses the first held-out discordance with GCF; a panel marked "Counterexample" instead shows GCF safe and the proposed policy unsafe. Selection is by scenario/seed order, not violation magnitude. Safe-trajectory counts for the full shock sets accompany the held-out results.

6. Discussion

6.1. State-dependent reserve retention near operating boundaries

The distinctive role of the filter is to regulate reserve consumption, not merely the next state's distance from a fixed boundary. Rearranging Equation (13) gives $h_i(x_{k+1}) - h_i(x_k) \geq -(1-\rho)h_i(x_k) + d_i$ under the nominal model. As the margin decreases, less of it can be used in one sample. Once it falls below $d_i/(1-\rho)$, the next nominal margin must increase. The intervals in Section 3.5 determine whether either device would need power outside its admissible range to achieve that increase.

The S2 results identify a local benefit of this mechanism. The uncertainty-aware policy completes 488/500 sampled-safe trajectories, compared with 470/500 for matched projection and 471/500 for nominal DT-CBF. Seventeen paired trajectories are safe with the uncertainty-aware policy but unsafe with nominal DT-CBF, with no reverse cases. The mechanism is therefore most relevant to boundary-proximal transients for which power redistribution or reserve recovery remains feasible. Away from the boundary in S1, matched hard projection and the uncertainty-aware filter behaved alike. Under the severe and long shocks in the additional families, the same offsets were insufficient to prevent all sampled violations. Even in S2, the safety gain was not accompanied by a shorter mean dwell time.

6.2. Relation to hard projection and guarded filtering

Matched hard projection and uncertainty-aware DT-CBF use the same nominal QP, operating set, input limits, solver and fallback on the same shock realizations. Only the constraint on the next margin changes. Nominal DT-CBF retains the proportional term alone, and the offset-only filter retains the allowance alone. In S2, these choices changed how many trajectories stayed within the sampled limits but had little effect on battery current.

GCF met the sampled safety limits on every S2 trajectory and left less energy unserved than the uncertainty-aware controller. Its battery carried more current. Mean RMS battery current was 28.28 A with the uncertainty-aware controller and 65.10 A with GCF. The similar currents among the QP-based controllers indicate that most of this difference came from the nominal power split, not from the barrier condition. The reconstructed current and the corresponding ohmic energy quantify electrical loading. Neither measures temperature rise or battery aging.

6.3. Applicability and model scope

Proposition 1 concerns the nominal energy model under commands that meet the filter constraints. For the sampled plant results in S1–S4, we used the synthetic component values in Table 1 and simulated losses, actuator dynamics and measurement errors. A physical installation would require identified component parameters and disturbance bounds compatible with its recovery power. Battery SOC changes slowly enough that a fixed offset can demand more inward motion in one sample than the battery can deliver, even with SOC well inside its physical limits.

The shock scenarios assume sustained missing supervisory feedforward: the shock remains absent from the current preview entry after onset. In the APCBF comparison (Section 5.5), providing the present bus demand but not the future shock changed which filter had the smaller mean absolute mismatch energy in S2. Recent HESS methods are compared by technique and control layer, and one recent

one-step constraint is adapted numerically, alongside the same-plant APCBF comparison. Full same-plant reproduction of the published Koopman, genetic-algorithm-tuned MPC and robust multivariable tracking controllers is outside this same-plant comparison. Local desktop timing establishes computational feasibility on the measured host, not an embedded worst-case execution bound. Experimental or hardware-in-the-loop validation may reveal operating-point-dependent ESR and polarization, converter tracking and communication delays, and SOC or SC-energy estimation errors beyond the fixed mismatch ranges considered here. Where credible bounds can be established, these measurements could inform updated model parameters, operating margins and barrier offsets. Temperature- or current-dependent limits would require additional state or input constraints. Such tests would show whether the boundary-proximal safety–service trade-off persists with measured terminal-voltage transients and real-time actuation.

7. Conclusions

This paper formulates an uncertainty-aware DT-CBF safety filter around a convex energy-domain HESS allocator. The filter limits depletion of the currently available SOC and SC-energy margins rather than constraining only the next state. This mechanism is intended for boundary-proximal operation in which demand can still be redistributed and the reserve can be recovered.

When an unpreviewed shock drove the system close to an operating boundary, more trajectories met the sampled physical limits with the proposed filter than with matched hard projection or nominal DT-CBF. Mean boundary dwell time did not decrease. GCF kept more trajectories within the limits and left less energy unserved, but required higher RMS battery current. The lower current with the proposed controller came mainly from the nominal power allocation. The filter determined how much reserve could be used at each step. In the additional shock and predictive-filter comparisons, which controller best tracked bus demand depended on pulse duration, initial reserve and the demand information available.

Adding the fixed offsets to nominal DT-CBF brought 17 more of the 500 paired S2 trajectories within the sampled limits. Across the 400 additional shocks, the two filters agreed on every trajectory. For a specific installation, the battery direction and both device power limits determine whether a chosen retention rate leaves enough power for recovery in one sample. The recovery intervals provide that check, and the offsets also need application-specific bounds on model mismatch. The nominal sampled-set result holds when a feasible filter command is applied at each step from an initially admissible state.

Data availability statement

No supplementary or additional data were generated in this study.

Declaration of generative AI and AI-assisted technologies

The authors used the generative artificial intelligence (AI) assistant OpenAI (model: GPT-6) for code optimization, language polishing. Figure 1 was created in Gemini Image. The authors did not use generative AI or AI-assisted technologies in the writing of this manuscript.The authors reviewed and edited the suggested changes, checked the implementation and numerical outputs, and take responsibility for the mathematical formulation, results and conclusions.

Acknowledgments

This work was supported in part by the National Natural Science Foundation of China under Grant 52307056.

Authors' contribution

Conceptualization, W.H. and W.Y.; methodology, W.H., Y.Z., and W.Y.; software, W.H.; validation, W.H. and Y.Z.; formal analysis, W.H. and Y.Z.; data curation, W.H.; visualization, W.H.; supervision and project administration, W.Y.; funding acquisition, W.Y.; writing—original draft preparation, W.H.; writing—review and editing, Y.Z. and W.Y. All authors have read and agreed to the published version of the manuscript.

Conflicts of interest

The authors declare no competing financial interests or personal relationships that could have influenced the work reported in this paper.

References

[1] Dougal RA, Liu S, White RE. Power and life extension of battery–ultracapacitor hybrids. *IEEE Trans. Compon. Packag. Technol.* 2002, 25(1):120–131.

[2] Lu S, Corzine KA, Ferdowsi M. A new battery/ultracapacitor energy storage system design and its motor drive integration for hybrid electric vehicles. *IEEE Trans. Veh. Technol.* 2007, 56(4):1516–1523.

[3] Cao J, Emadi A. A new battery/ultracapacitor hybrid energy storage system for electric, hybrid, and plug-in hybrid electric vehicles. *IEEE Trans. Power Electron.* 2012, 27(1):122–132.

[4] Song Z, Li J, Hou J, Hofmann H, Ouyang M, *et al.* The battery-supercapacitor hybrid energy storage system in electric vehicle applications: a case study. *Energy* 2018, 154:433–441.

[5] Kouchachvili L, Yaïci W, Entchev E. Hybrid battery/supercapacitor energy storage system for the electric vehicles. *J. Power Sources* 2018, 374:237–248.

[6] Babu TS, Vasudevan KR, Ramachandaramurthy VK, Sani SB, Chemud S, *et al.* A comprehensive review of hybrid energy storage systems: converter topologies, control strategies and future prospects. *IEEE Access* 2020, 8:148702–148721.

[7] Zhang Q, Deng W, Zhang S, Wu J. A rule based energy management system of experimental battery/supercapacitor hybrid energy storage system for electric vehicles. *J. Control Sci. Eng.* 2016, 2016:6828269.

[8] Mesbahi T, Rizoug N, Bartholomeüs P, Sadoun R, Khenfri F, *et al.* Optimal energy management for a Li-ion battery/supercapacitor hybrid energy storage system based on particle swarm optimization incorporating Nelder–Mead simplex approach. *IEEE Trans. Intell. Veh.* 2017, 2(2):99–110.

[9] Zhang S, Xiong R, Sun F. Model predictive control for power management in a plug-in hybrid electric vehicle with a hybrid energy storage system. *Appl. Energy* 2017, 185:1654–1662.

[10] Huang Y, Wang H, Khajepour A, He H, Ji J. Model predictive control power management strategies for HEVs: a review. *J. Power Sources* 2017, 341:91–106.

[11] Ma B, Guo X, Li P. Adaptive energy management strategy based on a model predictive control with real-time tuning weight for hybrid energy storage system. *Energy* 2023, 283:129128.

[12] Zhao Y, Qin F, Zhang Z. Robust tracking control design of hybrid battery–supercapacitor energy storage system in electric vehicles. *IEEE Trans. Transp. Electrif.* 2025, 11(1):3262–3276.

[13] Chen B, Wang M, Hu L, He G, Yan H, *et al.* Data-driven Koopman model predictive control for hybrid energy storage system of electric vehicles under vehicle-following scenarios. *Appl. Energy* 2024, 365:123218.

[14] Ma B, Li PH. Optimal flexible power allocation energy management strategy for hybrid energy storage system with genetic algorithm based model predictive control. *Energy* 2025, 324:135958.

[15] Hu Q, Xie S, Zhang J. Data-based power management control for battery supercapacitor hybrid energy storage system in solar DC-microgrid. *Sci. Rep.* 2024, 14:26164.

[16] Kumaresan N, Rammohan A. Fuzzy reinforcement learning-based energy management strategy for the integration of hybrid energy storage system of electric vehicle. *Results Eng.* 2025, 28:107875.

[17] Díaz-González F, Heredero-Peris D, Borrego-Orpinell G, Capó-Lliteras M. Modeling and impact assessment of hybrid battery–supercapacitor energy storage solutions for electric vehicles. *J. Energy Storage* 2026, 152:120811.

[18] Raju CG, Senapati L, Behera PK, Lenka RK, Ray P. Modeling, control, and experimental validation of a single-stage triple-port interfacing converter for photovoltaic-hybrid energy storage integration in direct current microgrid. *Energy Storage* 2026, 8(4):e70404.

[19] Mayne DQ, Rawlings JB, Rao CV, Scokaert POM. Constrained model predictive control: stability and optimality. *Automatica* 2000, 36(6):789–814.

[20] Mayne DQ, Seron MM, Raković SV. Robust model predictive control of constrained linear systems with bounded disturbances. *Automatica* 2005, 41(2):219–224.

[21] Ames AD, Xu X, Grizzle JW, Tabuada P. Control barrier function based quadratic programs for safety critical systems. *IEEE Trans. Autom. Control* 2017, 62(8):3861–3876.

[22] Ames AD, Coogan S, Egerstedt M, Notomista G, Sreenath K, *et al.* Control barrier functions: theory and applications. In *Proceedings of the European Control Conference*, Naples, Italy, June 25–28, 2019, pp. 3420–3431.

[23] Agrawal A, Sreenath K. Discrete control barrier functions for safety-critical control of discrete systems with application to bipedal robot navigation. In *Proceedings of Robotics: Science and Systems XIII*, Cambridge, USA, July 12–16, 2017.

[24] Zeng J, Zhang B, Sreenath K. Safety-critical model predictive control with discrete-time control barrier function. In *Proceedings of the American Control Conference*, New Orleans, USA, May 25–28, 2021, pp. 3882–3889.

[25] Zeng J, Li Z, Sreenath K. Enhancing feasibility and safety of nonlinear model predictive control with discrete-time control barrier functions. In *Proceedings of the 60th IEEE Conference on Decision and Control*, Austin, USA, December 13–15, 2021, pp. 6137–6144.

[26] Khajenejad M, Cavorsi M, Niu R, Shen Q, Yong SZ. Tractable compositions of discrete-time control barrier functions with application to driving safety control. In *Proceedings of the European Control Conference*, Online, June 29–July 2, 2021, pp. 1303–1309.

[27] Xiong Y, Zhai DH, Tavakoli M, Xia Y. Discrete-time control barrier function: high-order case and adaptive case. *IEEE Trans. Cybern.* 2023, 53(5):3231–3239.

[28] Ahmed M, Masood U, Azeem MK, Ahmad I, Jabbar AU. Barrier function based adaptive sliding mode controller for the hybrid energy storage system of plugin hybrid electric vehicles. *J. Energy Storage* 2023, 72:108051.

[29] Shaw Cortez W, Drgona J, Vrabie D, Halappanavar M. A robust, efficient predictive safety filter. *arXiv* 2024, arXiv:2311.08496v2.

[30] Didier A, Zeilinger MN. Approximate predictive control barrier function for discrete-time systems. *IEEE Trans. Autom. Control* 2026, 71(7):4905–4912.

[31] Stellato B, Banjac G, Goulart P, Bemporad A, Boyd S. OSQP: an operator splitting solver for quadratic programs. *Math. Program. Comput.* 2020, 12:637–672.