ARFLO RESEARCH / TECHNICAL PAPER 004

A physics-based filter for cooling commands.

arflo GmbH

4 November 2026 / Edition 1.0

Abstract

Learning and optimization can propose cooling actions without explicitly enforcing plant constraints. This paper derives a minimal-intervention filter for a scalar thermal energy balance with bounded heat input and bounded temperature-estimation error. Exact integration gives the minimum delivered heat-removal rate needed over the next control interval. A numerical example checks the command, its thermal units and its feasibility against actuator capacity. Intersample temperature behavior and the conditions for repeated feasibility are stated explicitly. The analysis is conditional on the model and input bounds; it does not establish installation-wide safety or functional-safety certification.

1 Separate the proposed command from its admissibility.

A cooling controller may optimize electrical demand, tracking error or another objective. Its preferred action still needs to fit temperature and actuator constraints. Safety-filter research formalizes this separation: accept a proposed input when admissible, otherwise replace it with a constraint-respecting alternative [1, 2]. The derivation here specializes that architectural idea to a transparent, single-state thermal model; it is not a reproduction of the general predictive-filter theorem.

The constrained temperature must correspond to a specified physical location. A coolant-return limit, an air-inlet limit and a semiconductor-junction limit are different quantities. The permitted value must come from the actual equipment and operating agreement. No generic numerical temperature limit is assigned to a hardware family in this paper.

2 A bounded thermal energy balance.

Let θ = T - T_in be temperature rise above a constant reference temperature. The model is C·dθ/dt = P - Gθ - u. C > 0 is thermal capacity in J/K; G > 0 is passive conductance in W/K; P is heat input in W; and u is delivered active heat removal in W. Delivered heat removal is a thermal quantity, distinct from electrical pump or chiller power. A pump-speed or valve command needs a validated plant map before it can be used as u.

Assume known fixed C and G, P(t) ≤ P_bar over the next interval, θ at its start no greater than θ_bar, and immediate delivery of a constant removal u. θ_bar may combine an estimate with a deterministic error bound. It is not automatically a Gaussian confidence interval. Comparing against the trajectory driven by P_bar yields an upper temperature envelope. These assumptions deliberately exclude actuator delay, unbounded faults and changing heat-transfer parameters.

$$\bar{\theta}^{+} = a\bar{\theta} + b(\bar{P} - u), \qquad a = e^{-G\Delta t/C}, \quad b = \frac{1-a}{G} \tag{1}$$

The coefficient a is dimensionless and b has units K/W. This is exact for the constant-input comparison model.

Research origin: Memo 022 (2026-04-04). Online edition: `https://arflo.io/research/physics-based-cooling-safety-filter/`.

3 Derive the smallest admissible intervention.

Requiring the upper prediction to stay below θ_lim gives an explicit lower bound u_req on heat removal. The physical actuator interval is $0 \leq$ u $\leq$ u_max. Define the admissible lower endpoint u_low = max(0, u_req). If u_low $\leq$ u_max, the closest feasible command to a nominal proposal u_nom is the projection of u_nom onto that interval. For this scalar quadratic objective no iterative optimization solver is needed.

If the lower endpoint exceeds available capacity, the filter has no admissible solution. Selecting maximum cooling may be a necessary response, but does not restore a mathematical guarantee. The operating design must specify another means of reducing the heat load or moving the system to a protected state. This infeasibility indication is an essential output of the filter.

$$u_{\mathrm{req}} = \bar{P} - \frac{\theta_{\mathrm{lim}} - a\bar{\theta}}{b} \tag{2}$$

This bound concerns heat removal over the specified interval, not an electrical power-saving target.

$$u^* = \min\{u_{\max}, \max(u_{\mathrm{nom}}, u_{\mathrm{low}})\}, \qquad u_{\mathrm{low}} \leq u_{\max} \tag{3}$$

The projection is used only after feasibility has been established.

4 Worked example with a capacity check.

Use C = 10,000 J/K, G = 20 W/K and Δt = 10 s. The time constant is 500 s. Choose θ_lim = 30 K, an estimated θ = 28.5 K and a deterministic error bound of 0.3 K, giving θ_bar = 28.8 K. Let P_bar = 2,000 W, u_max = 1,500 W and the nominal removal proposal be 100 W. These are illustrative model parameters, not a measured installation or a device rating.

The exact coefficients are a = 0.9801986733 and b = 0.0009900663 K/W. The required removal is 211.960 W. The 100 W proposal would allow the upper prediction to reach 30.1108 K, so it is replaced by 211.960 W. At that filtered input the envelope reaches exactly 30 K. Numerical implementation should preserve a specified positive margin for rounding and solver tolerance, with its value included in the design.

Table 1: Upper-envelope temperature rise after the 10 s interval.

Delivered removal / W	**Predicted rise / K**	**Interpretation**
100.000	30.1108	Nominal proposal
211.960	30.0000	Filtered boundary
1500.000	28.7248	Maximum removal

5 Between samples and over repeated intervals.

An endpoint constraint alone is generally insufficient for a multi-state or delayed plant. In this particular scalar comparison model, constant input produces a monotonic approach to the equilibrium (P_bar - u)/G. Its maximum on a closed interval is therefore at an endpoint. A safe starting upper bound and a safe final upper prediction imply that the entire comparison trajectory is below the limit. Since actual heat input is no greater than P_bar, the comparison also bounds the actual trajectory under the stated model.

Repeated feasibility additionally needs enough cooling authority at the limit. For fixed bounds, immediate actuation and no rate limit, u_max $\geq$ P_bar - Gθ_lim is a sufficient capacity condition

for the scalar safe half-line to be maintainable. In the example the required boundary removal is 1,400 W, below the available 1,500 W. This verifies the ideal model's capacity condition; a changing load bound or actuator fault can invalidate it.

6 Delay, model mismatch and validation.

When a command takes effect only after a delay, first propagate through that delay with the previously delivered removal and worst-case heat input. Check the intermediate trajectory before computing the next command. Finite ramp rates require an actuator state or an explicit reachable-removal bound. Uncertain conductance, heat capacity, inlet temperature or delivered removal must likewise be included in a justified envelope; substituting their nominal values does not preserve robustness.

Validation should test infeasible load steps, sensor underestimation, delayed commands, communication loss and cooling-capacity degradation. Measure delivered heat removal and the constrained temperature at a resolution sufficient to detect intersample peaks. Lower-temperature constraints, coupled hotspots and hydraulic interactions require a richer model. A scalar filter is a useful auditable component, but certification and installation protection require evidence about the complete sensing, computation and actuation chain.

Conclusion

A cooling command is admissible only relative to a stated model, uncertainty envelope and available heat-removal capacity. An empty feasible set must remain visible.

References

[1] A. D. Ames, X. Xu, J. W. Grizzle and P. Tabuada (2017), Control Barrier Function Based Quadratic Programs for Safety Critical Systems, IEEE Transactions on Automatic Control 62(8), 3861-3876.. `https://arxiv.org/abs/1609.06408`.

[2] K. P. Wabersich and M. N. Zeilinger (2021), A predictive safety filter for learning-based control of constrained nonlinear dynamical systems, Automatica 129, 109597.. `https://arxiv.org/abs/1812.05506`.