Research Square

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

IT-power feedforward halves pump energy in liquid-cooled data centres

Yang Liu

1620872416@qq.com

Independent researcher

Article

Keywords:

Posted Date: October 29th, 2026

DOI: https://doi.org/10.21203/rs.3.rs-11480780/v1

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

Additional Declarations: No competing interests reported.

Abstract

Pump work in liquid-cooled data centres scales with the cube of flow rate, so the choice of controlled variable sets the achievable energy floor. Existing pump controllers regulate temperature — the result of the thermal balance — although the heat load (IT power), its cause, is directly readable from power distribution units. Here we show that replacing temperature feedback with a closed-form feedforward from IT power reduces pump work by 56.2% and peak junction temperature by 5.14 °C, against a temperature controller tuned over 90 parameter sets, in a 90 kW rack under a 30%→100% step. Conversely, a 300 s look-ahead increases pump work by 45.5% without changing peak temperature, because required flow is an instantaneous function of power. Inter-rack reallocation is likewise infeasible: each lower bound is set by its own load and cannot be transferred. The gain lies in the observed variable, not in algorithmic sophistication.

Introduction

Cooling is the second-largest electricity load in a data centre, and as racks move from air to liquid cooling the pump that circulates coolant has become the largest single consumer inside the cooling loop. Almost every pump controller in service regulates temperature: it waits for the coolant to warm before increasing flow, because temperature is the quantity that must not exceed its limit. Here we show that this choice of controlled variable, rather than any shortcoming of the control algorithm, sets a floor on pump energy that tuning cannot breach. The heat load is already measured at the power distribution unit, and it maps to the required flow through a closed-form relation, so the cause of the thermal excursion is available before the excursion occurs. The saving that follows is large, and it comes from what the controller observes rather than from how clever the controller is.

Rack power densities in artificial-intelligence clusters now reach 50–120 kW, a range in which air cooling is no longer viable and direct-to-chip liquid cooling has become the default thermal architecture[1,2]. Global data-centre electricity demand is projected to roughly double between 2024 and 2030, driven largely by such clusters[3], so the auxiliary power consumed by the cooling loop itself has become a first-order design target rather than a secondary concern[4,5].

For the secondary (facility-water) loop, pump power follows the affinity laws[6,7] and therefore grows with the cube of flow rate. Because flow is set by the controller, the *choice of what the controller observes* sets the energy floor that no tuning can go below. Three families of pump controllers dominate practice: rule-based schemes that step pump speed on temperature thresholds; model-predictive control, which optimises over a receding horizon using a thermal or hydraulic model; and learning-based schemes that forecast load or temperature from historical data. All three regulate temperature.

This choice is historically natural — temperature is the constraint that must ultimately be satisfied — but it neglects an asymmetry specific to liquid-cooled facilities: **IT power is measurable in real time, with negligible latency, at the power distribution unit or server baseboard management controller, and it**

relates to the required flow through a purely algebraic relation. The cause of the thermal excursion is available; the control loop nevertheless consumes the effect.

Here we quantify what that substitution is worth, and test two optimisation strategies that are widely assumed to help. Our contributions are: (i) a closed-form flow demand and a five-statement feedforward control law requiring no PID loop, no forecasting model and no learning, which reduces pump work by 56.2% against a parameter-swept temperature controller; (ii) a demonstration that look-ahead flow elevation is a *negative* return under measurable IT power, with a stated condition for when it becomes useful; and (iii) a demonstration that inter-rack flow reallocation cannot save energy because the lower bounds are not transferable resources.

Results

Physical basis and constraint set

The model rests on three relations and no others: energy conservation, $\Delta T = P_H/(\rho\, c_p\, Q)$; the quadratic network characteristic, $H = S\, Q^2$; and the resulting pump power, $P_p = \rho\, g\, S\, Q^3/\eta$. The cubic form holds **only when pump speed alone is varied and valve positions remain fixed**, since S is a function of valve opening; the present work therefore restricts the actuator to a variable-frequency drive, with valves used solely for initial hydraulic balancing.

Three hard constraints apply: junction temperature $T_{chip} \le 85$ °C (S1)[8]; supply–return temperature difference $\Delta T \le 15$ °C (S2)[9,10]; and actuator slew rate $|dQ/dt| \le$ RATE (S3). The control objective is to minimise $\int P_p\, dt$ subject to S1–S3, with no weighting or trade-off permitted among them.

Both S1 and S2 can be inverted to give closed-form lower bounds on flow. For S1, $T_{chip} = T_{in} + P_H R_{th} + P_H/(2\rho\, c_p\, Q) \le 85$ gives $Q \ge P_H/[2\rho\, c_p(85 - T_{in} - P_H R_{th})]$; for S2, $Q \ge P_H/(\rho\, c_p \cdot 15)$. The required flow is the larger of the two. **For an 80 kW rack under the present parameters, the S1 bound is 31.9 L min⁻¹ and the S2 bound 76.6 L min⁻¹, so the binding constraint is temperature difference rather than junction temperature.** This distribution governs the interpretation of the reallocation result below.

Junction temperature does not respond instantaneously; a first-order lumped model gives $C_{chip}\, dT_{chip}/dt = P_H - (T_{chip} - T_w)/R_{th}$ with $T_w = T_{in} + P_H/(2\rho\, c_p\, Q)$. The thermal time constant $\tau = R_{th} C_{chip}$ is 20 s for the present parameters, whereas the actuator needs 60 s to traverse its full range, so the actuator is three times slower than the thermal inertia.

Feedforward versus a parameter-swept temperature controller

The control law comprises five statements: read IT power; evaluate the closed-form flow demand; clip the command to the slew-rate envelope; clip to pump limits; map to drive frequency via $Q \propto n$. It is

constant-time (sub-microsecond) and contains no feedback path — temperature is monitored for over-temperature protection but does not enter the steady-state flow calculation.

A 90 kW rack was simulated for 1800 s under a load profile of 30% (0–600 s), a step to 100% (600–1200 s), and 40% (1200–1800 s). The temperature controller $q_{cmd} = Q + k_p(T_{chip} - T_{sp})$ was tuned over the full grid $T_{sp} \in \{40, 42, 45, 47, 50, 52, 55, 58, 60, 65\}$ °C × $k_p \in \{2\times10^{-5} \ldots 1\times10^{-3}\}$, 90 combinations, all of which satisfied the constraints; the lowest-energy member (T_{sp} = 65 °C, $k_p = 5\times10^{-5}$) was retained as the baseline.

Feedforward consumed 0.0765 kWh against 0.1747 kWh for the tuned controller, a **reduction of 56.2%** (Fig. 1b,c), with peak junction temperatures of 78.50 °C and 83.64 °C respectively — **5.14 °C lower** (Table 1). The feedforward flow ranged over 26–86 L min⁻¹ versus 14–117 L min⁻¹ for the feedback loop: the feedback controller both under-supplies during the transient and over-supplies in steady state (Fig. 1b).

To exclude the possibility that the baseline was weakened, the set point was varied independently at fixed $k_p = 1\times10^{-4}$. Pump work fell monotonically with increasing set point — 0.5286, 0.3636, 0.1943, 0.1787 and 0.1747 kWh at 40, 50, 55, 60 and 65 °C — and saturated near 65 °C. The reported 56.2% is therefore measured against a saturated optimum, not against an arbitrary set point.

Look-ahead flow elevation as a negative return

A common proposal is to raise pump speed before a load increase. Implemented in its safe form — the commanded flow bounded below by the flow demand of max(P_now, P_future) — look-ahead cannot violate current constraints, but its benefit remains to be shown.

Across look-ahead windows of 0, 60, 120, 300 and 600 s, pump work rose monotonically from 0.0765 to 0.0835, 0.0904, 0.1113 and 0.1469 kWh, while **peak junction temperature remained 78.50 °C in every case** (Fig. 2). At 300 s the penalty is 45.5% with no thermal benefit whatsoever.

Two reasons account for this. First, the flow demand is an instantaneous function of power; before power rises, added flow serves no constraint, since the temperature is already within limits. Second, although the actuator is three times slower than the thermal inertia, the flow increment required for a 30%→100% step is a small fraction of the pump range, so a 60 s ramp suffices and a five-minute lead is superfluous. Look-ahead acquires value only when the actuator response greatly exceeds τ *and* the step approaches full range; in that regime, relaxing RATE by specifying a faster drive or increasing C_{chip} by thermal storage dominates any control-side measure. We note the scoping condition explicitly: the look-ahead tested here uses *current* power. If future load is available from a scheduling layer, look-ahead carries genuine information and this negative result does not apply.

Inter-rack reallocation is not a saving mechanism

For three racks of 80, 50 and 30 kW sharing a common header ($S_{main} = 1.2\times10^6$, branch $S = 3.0\times10^6$ each), the individual demand bounds are 76.6, 47.8 and 28.7 L min⁻¹, and operating at these bounds

consumes 0.4582 kW.

Reducing one rack's flow by 5%, 10% or 20% and adding the same quantity to another lowers total pump power by 1.76%, 2.98% and 3.79% respectively — **but each case immediately violates the donor rack's 15 °C temperature-difference constraint.** Savings of this kind are purchased by breaking a constraint and are therefore not admissible (Fig. 3). Conversely, if all branches are oversized by 50% the system consumes 1.5463 kW; returning to the bounds saves 70.4% (Fig. 3).

The underlying reason is structural rather than control-theoretic: each rack's lower bound is set by its own thermal load, and lower bounds are not exchangeable resources. The energy resides in withdrawing oversizing, not in moving flow between branches.

Discussion

The two controllers differ not in sophistication — the proposed law has five statements and no optimiser — but in the variable observed. Temperature is the result of the thermal balance: it must first move before the loop can respond, and the resulting integral error scales with the ratio of actuator response to thermal time constant. IT power is the cause: its effect on flow demand follows algebraically, without an error–integrate–correct path. This explains why feedforward both saves energy and lowers the peak: it does not execute the same correction faster but prevents the deviation from arising.

The reallocation result should not be read as a claim that multi-rack coordination is pointless. The correct statement is narrower: when every branch must independently satisfy its own constraint, the minimum-energy allocation is the vector of bounds, and no transferable margin exists. The binding constraint in the present parameter range is the 15 °C temperature difference, not the 85 °C junction limit. Temperature-difference specification is therefore a first-order operating parameter rather than a matter of convention — a point consistent with measurements showing that supply-water temperature materially shifts whole-facility energy in direct liquid-cooled installations[2].

We also record a negative result concerning topology-based diagnostics. Computing the homology of a pipe-network complex and flagging changes in Betti number is blind to the dominant failure mode: when a branch is partially blocked or its valve mis-set and flow departs by 20% from its design share, the network remains connected and H_0 is unchanged. Only genuine disconnection alters H_0, and such events are a minority. Conservation-based residuals — nodal mass-balance error, share-based z-scores, and cross-checks between temperature difference and flow that discriminate a real blockage from a faulty sensor — are sensitive to precisely the phenomena topology ignores.

Several limitations bound these findings. All results are numerical, from a one-dimensional lumped-parameter model; they indicate magnitude and trend, not site-specific performance. Multi-chip non-uniform heat flux, manifold imbalance, and along-path thermal stratification are not represented. No hardware was interfaced: the output is a pump-speed command, without drive communication, failover, redundancy interlocking, or sensor-fault detection. Thermal storage as an alternative peak-absorption

route was not costed. Multi-level manifold networks introduce stronger hydraulic coupling than the single-header topology analysed here, and deserve separate treatment.

Methods

System parameters

Table 1 | Parameters used in all simulations.

Parameter	Value
Rated thermal load	90 kW
Supply water temperature T_{in}	35 °C
Chip-to-coldplate thermal resistance R_{th}	4×10^{-4} K/W
Equivalent chip thermal capacitance C_{chip}	5×10^{4} J/K
Pump rating	90 L min^{-1} at 20 m head, $\eta = 0.6$
Network resistance coefficient S	8.89×10^{6} (from $H = S Q^2$ at rated point)
Slew rate RATE	full range / 60 s
Pump range	0.15–1.3 × rated flow (13.5–117 L min^{-1})
Water density / volumetric heat capacity $\rho\, c_p$	1000 kg m^{-3} / 4.18×10^{6} J m^{-3} K^{-1}
Gravitational acceleration g	9.81 m s^{-2}
Simulation step Δt	1 s
Simulation duration	1800 s

Load profile

$P_H(t) = 0.30\, P_{rated}$ for $0 \le t < 600$ s; $1.00\, P_{rated}$ for $600 \le t < 1200$ s; $0.40\, P_{rated}$ for $1200 \le t < 1800$ s. The 600 s transition is an instantaneous step, representing the onset of a training workload.

State integration

At each step: compute $P_H(t)$; evaluate the commanded flow; clip to $[Q_{prev} - RATE\cdot\Delta t, Q_{prev} + RATE\cdot\Delta t]$; clip to $[Q_{min}, Q_{max}]$; set Q; compute mean coldplate water temperature $T_w = T_{in} + P_H/(2\rho\, c_p\, Q)$; integrate T_{chip} by explicit Euler with the stated Δt; accumulate $P_p\cdot\Delta t$; track the running maximum of T_{chip}.

Flow demand

$q_{need}(P) = \max(q_{S1}(P), q_{S2}(P))$
$q_{S1}(P) = P / [2 \rho c_p (85 - T_{in} - P R_{th})]$
$q_{S2}(P) = P / (\rho c_p \cdot 15)$
safeguarded: the q_{S1} denominator is floored at 1e-3

Both are closed-form; no iteration is performed.

Pump power

$P_p = \rho g S Q^3 / \eta$, evaluated in kW as $\rho g S Q^3 / (1000 \eta)$.

Controllers compared

1. **Temperature feedback (baseline).** $q_{cmd} = Q + k_p (T_{chip} - T_{sp})$, swept over the 90-point grid given in Results; the minimum-energy constraint-satisfying member is reported.
2. **Feedforward.** $q_{cmd} = q_{need}(P_H)$.
3. **Feedforward with look-ahead.** q_{cmd} = max over $k \in \{0, L/2, L\}$ of $q_{need}(P_H(t+k))$, for L = 60, 120, 300, 600 s.
4. **Naive look-ahead (reported for safety only).** $q_{cmd} = q_{need}(P_H(t+L))$ without a lower bound at the current demand; this violates S1 and is included to quantify the hazard.

Multi-rack case

Three racks of 80, 50, 30 kW; a common header of $S_{main} = 1.2 \times 10^6$ and three branches of $S = 3.0 \times 10^6$. Total pump power is $\rho g [S_{main} (\Sigma Q_i)^3 + \Sigma S_i Q_i^3]/\eta$. Reallocation cases move a fraction $f \in \{5, 10, 20\}$% of one branch's bound flow to another and record both the power change and whether the donor's S2 bound is still satisfied.

Reproducibility

The model contains no random numbers; a given parameter set reproduces identical output bit-for-bit on the same implementation. All quantities in Tables and text derive from the single implementation described above; no measured data were used.

Declarations

Data availability

All data generated in this study are contained in the Results and Methods sections, and are reproducible from the model described therein. No experimental datasets were generated or analysed.

Code availability

The complete implementation — model, controllers, parameter sweep and multi-rack case — is available at https://gitee.com/liu-yang1215l/tft-lnn-pipeline under projects/liquid_cooling/, comprising cooling_opt.py (model and initial comparison), final_compare.py (parameter sweep and look-ahead scan), and alloc_test.py (multi-rack reallocation). The code requires only NumPy; no proprietary solver is used.

References

1. Zhang, Y., Zhao, Y., Dai, S. et al. Cooling technologies for data centres and telecommunication base stations — a comprehensive review. *J. Clean. Prod.* **334**, 130280 (2022).
2. Gheni, M., Kerskes, H. & Stergiaropoulos, K. Operational analysis of the cooling system in a direct liquid-cooled data center: a measurement and simulation study on the impact of supply water temperature. *Appl. Energy* (in press).
3. International Energy Agency. *Energy and AI* (IEA, Paris, 2025).
4. He, Z., Xi, H., Wang, J. et al. Synergy optimization analysis of heat transfer performance and energy consumption in heat transfer process and its application in data centers. *Appl. Energy* **307**, 118276 (2022).
5. Enhancing data center cooling efficiency and ability: a comprehensive review of direct liquid cooling technologies. *Energy* (2024).
6. Karassik, I. J., Messina, J. P., Cooper, P. & Heald, C. C. *Pump Handbook* 3rd edn (McGraw-Hill, 2001).
7. Open Compute Project. *Advanced Cooling Facility Guidelines* (OCP, 2021).
8. ASHRAE Technical Committee 9.9. *Thermal Guidelines for Data Processing Environments* 4th edn (ASHRAE, 2015).
9. Ministry of Housing and Urban-Rural Development of the People's Republic of China. Code for Design of Heating, Ventilation and Air Conditioning in Civil Buildings GB 50736—2012 (China Architecture & Building Press, 2012).
10. Lu Yaoqing. Practical Heating and Air Conditioning Design Manual, 3rd edn. (China Architecture & Building Press, 2020)

Figures

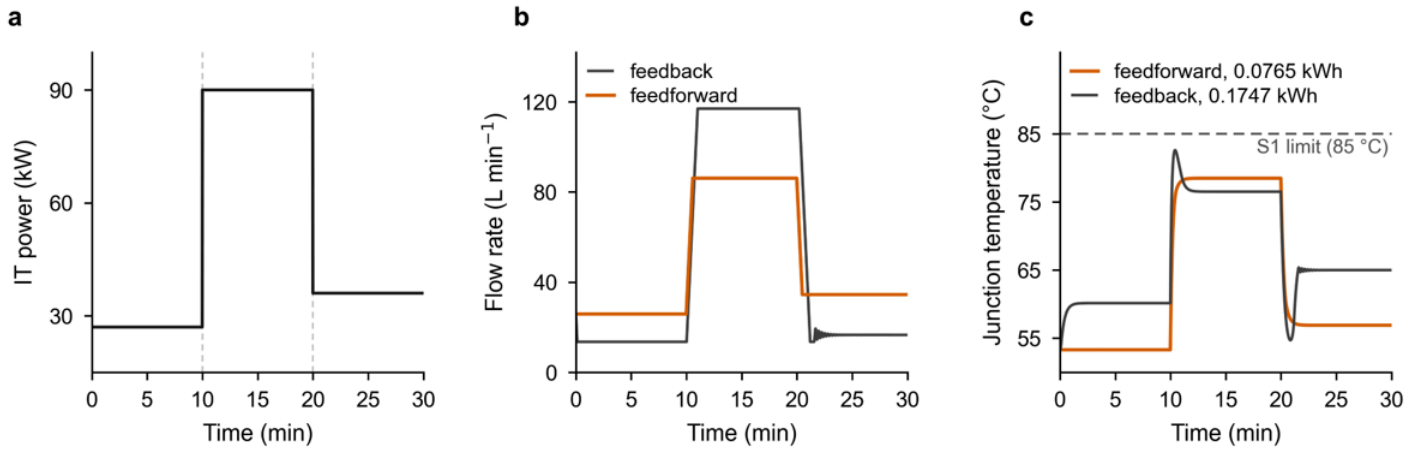


Figure 1

Feedforward from IT power versus a parameter-swept temperature controller. a, IT load profile (30% for 0–10 min, step to 100% for 10–20 min, 40% for 20–30 min; 90 kW rated). b, Commanded flow. The feedback loop (grey) over-supplies in steady state and lags the step; the feedforward law (red) tracks the closed-form demand. c, Junction temperature. Numbers in the legend give total pump work over the 30-min window. Dashed line, constraint S1 (85 °C).

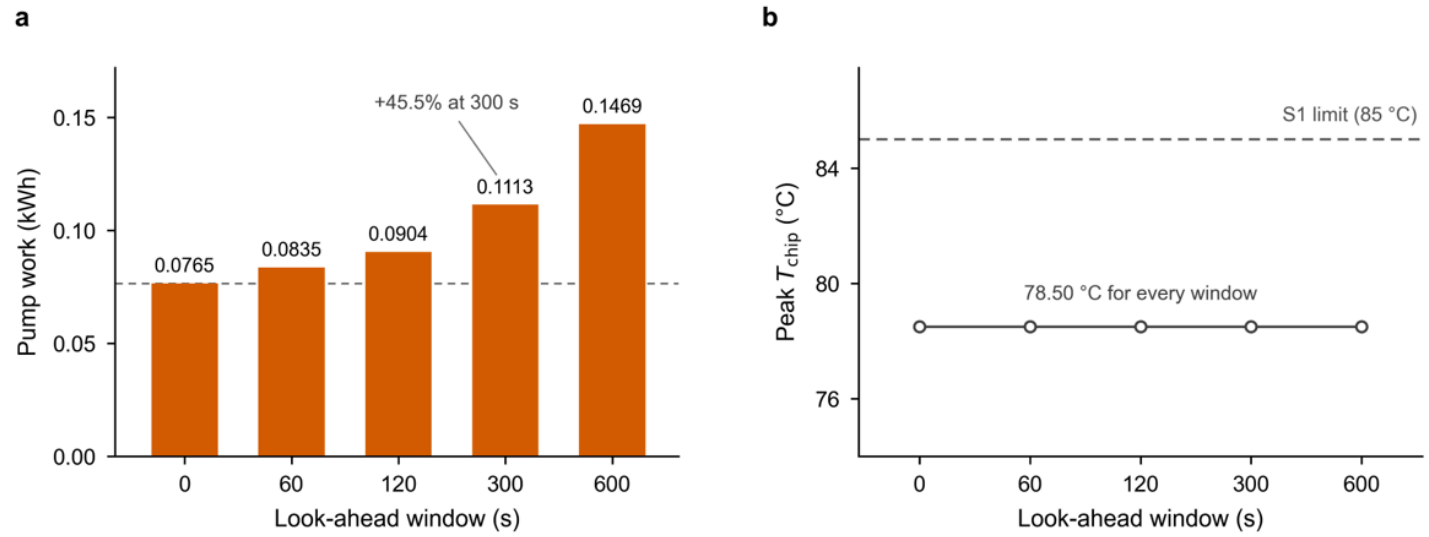


Figure 2

Look-ahead flow elevation is a negative return. a, Pump work against look-ahead window; monotonic increase from 0.0765 to 0.1469 kWh. b, Peak junction temperature, identical (78.50 °C) for every window. Dashed line, constraint S1.

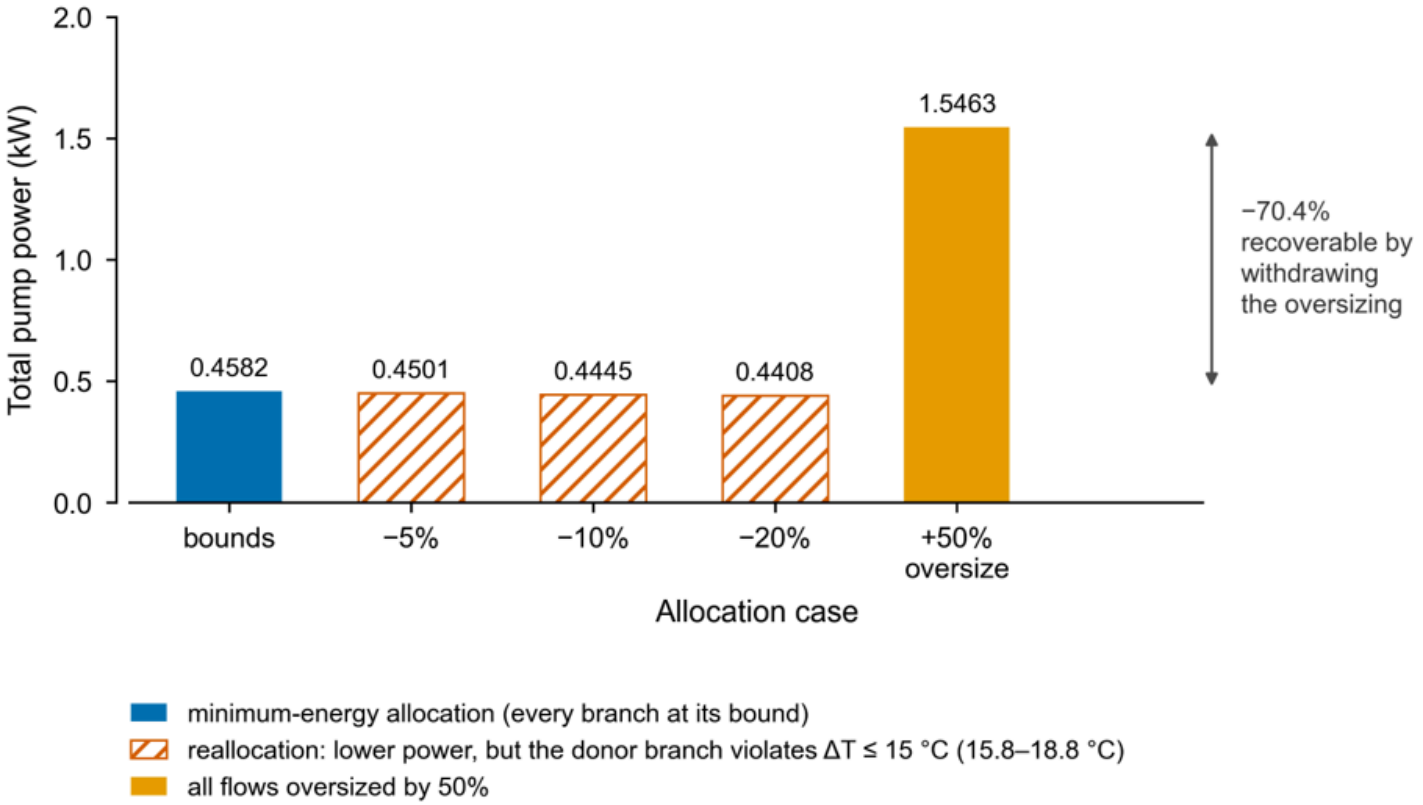


Figure 3

Inter-rack flow reallocation. Total pump power for three racks (80, 50, 30 kW) sharing a common header. Bars 2–4 move 5, 10 and 20% of one branch's bound flow to another; each lowers total power but violates the donor branch's ΔT ≤ 15 °C constraint. The final bar shows 50% oversizing, whose withdrawal — not reallocation — is the available saving.