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.

Benchmarking Classical, Hybrid, and Quantum Annealing Workflows on a Structured Control QUBO

Hamze Alavirad

h.alavirad@qusmos.com

QUSMOS

Maryam Bahrami Zanjani

QUSMOS

Research Article

Keywords: quantum annealing, quadratic unconstrained binary optimization, simulated annealing, path-integral annealing, hybrid quantum-classical optimization, optimization benchmarking

Posted Date: July 17th, 2026

DOI: https://doi.org/10.21203/rs.3.rs-10328048/v1

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

Additional Declarations: No competing interests reported.

Benchmarking Classical, Hybrid, and Quantum Annealing Workflows on a Structured Control QUBO

Hamzeh Alavirad[1*] and Maryam Bahrami Zanjani[1]

[1]QUSMOS, Karlsruhe, Germany.

*Corresponding author(s). E-mail(s): h.alavirad@qusmos.com;
Contributing authors: m.bahrami@qusmos.com;

Abstract

We benchmark classical, simulated-quantum, hybrid, and direct quantum-annealing workflows on a structured binary control problem formulated as a quadratic unconstrained binary optimization (QUBO) model. The benchmark is motivated by greenhouse heater scheduling, where the horizon $\boldsymbol{H}$ denotes the number of hourly binary control decisions, but the focus is on solver behavior for a physically decoded control-QUBO formulation. For the main one-day instance ($\boldsymbol{H} = \mathbf{24}$), all solver outputs are decoded back into heater schedules and evaluated using the same simulator-side objective and feasibility criterion. Classical simulated annealing and path-integral simulated quantum annealing produce feasible near-optimal solutions in all repetitions, with best objectives close to the exact optimum. In contrast, the tested D-Wave Leap Hybrid binary quadratic model workflow is less reliable and does not outperform the classical baselines under 15–60 s requested time limits. Direct D-Wave quantum processing unit execution on reduced instances remains feasible in all runs and recovers the exact optimum for $\boldsymbol{H} = \mathbf{10}$ and $\boldsymbol{H} = \mathbf{12}$, but the exact-hit rate decreases with increasing horizon. The results do not indicate quantum advantage, but provide a reproducible benchmark for evaluating annealing-based natural and quantum computing workflows on structured control QUBOs.

Keywords: quantum annealing, quadratic unconstrained binary optimization, simulated annealing, path-integral annealing, hybrid quantum-classical optimization, optimization benchmarking

MSC Classification: 90C27 , 90C59 , 81P68 , 93C55

1 Introduction

Natural and quantum-inspired computing methods provide alternative approaches for solving structured combinatorial optimization problems. In particular, annealing-based methods translate optimization into the search for low-energy configurations of an effective energy landscape, which makes them naturally compatible with binary quadratic formulations Jiang et al. (2023); D-Wave Systems Inc. (2026). Quadratic unconstrained binary optimization (QUBO) is a useful modelling interface because it can be addressed by classical heuristics, simulated quantum annealing methods, hybrid quantum-classical workflows, and quantum annealing hardware Glover et al. (2018); D-Wave Systems Inc. (2026a,b). Related gate-based approaches, such as the Quantum Approximate Optimization Algorithm, also motivate the use of binary optimization formulations as benchmark interfaces for near-term and future quantum optimization workflows Farhi et al. (2014).

However, for application-motivated control problems, raw QUBO energy is often not the final performance criterion. Candidate binary solutions must be decoded back into physically meaningful control trajectories and evaluated under the original simulator-side objective and feasibility constraints. This distinction is particularly important when benchmarking quantum and quantum-hybrid workflows, because embedding, penalty scaling, and solver-specific sampling behavior can affect the relationship between BQM energy and decoded physical performance D-Wave Systems Inc. (2026a,b); King et al. (2015).

This paper studies this issue through a structured binary control benchmark motivated by greenhouse heater scheduling. The greenhouse model is intentionally reduced; its role is to provide a transparent, physically decoded control-QUBO instance rather than a deployable greenhouse controller.

Greenhouse climate control is an important optimization problem because environmental regulation directly affects crop performance, resource consumption, and operating cost. In practical greenhouse operation, climate management depends on several strongly coupled variables, including temperature, humidity, radiation, and CO_2, together with external disturbances, actuator limitations, and economic constraints. This makes greenhouse climate control a complex decision-making problem in which model quality, optimization strategy, and computational tractability all play a central role Bersani et al. (2020); Morcego et al. (2023); Chen et al. (2025).

A broad range of methods has been investigated to optimize greenhouse climate control systems. Model-based approaches, in particular model predictive control (MPC), are widely used because they combine process dynamics, operational constraints, and economic or agronomic objectives within a receding-horizon framework Bersani et al. (2020); Morcego et al. (2023). In parallel, data-driven and learning-based approaches have been explored to improve prediction and decision-making under uncertainty Ajagekar et al. (2023); Mansour et al. (2025). Surrogate-based optimization has also been proposed to determine environmental targets from learned crop-response models. In particular, Lu et al. (2023) introduced an approach based on support vector regression and a quantum genetic algorithm, which should be

understood here as a quantum-inspired classical metaheuristic rather than direct execution on quantum hardware. Related earlier work used support-vector-machine-based models for greenhouse light-environment optimization Xin et al. (2019).

These studies show that greenhouse control has already moved beyond purely rule-based operation toward optimization-driven and learning-assisted formulations.

From the perspective of quantum computing, this class of problems is relevant because discretized control formulations can lead to large structured combinatorial optimization problems. After suitable simplification, binary encoding, and penalty-based reformulation, such problems can be expressed as quadratic unconstrained binary optimization (QUBO) models Glover et al. (2018); D-Wave Systems Inc. (2026). This is important because QUBO provides a common interface between the application-side control problem and several quantum-related optimization paradigms. In the near term, quantum-inspired optimization methods have shown promising behavior on hard combinatorial problems, suggesting that physics-inspired search strategies may be useful even on classical or specialized non-gate-based hardware Jiang et al. (2023). In parallel, quantum-native optimization approaches such as the Quantum Approximate Optimization Algorithm (QAOA) provide a longer-term route for solving binary optimization problems on gate-based quantum computers Farhi et al. (2014). Although large-scale fault-tolerant quantum optimization is not yet available in practice, these developments motivate expressing control-oriented benchmark problems in forms that are compatible with both classical and quantum optimization workflows.

Against this background, the present work does not aim to propose a new greenhouse controller or a new agronomic set-point strategy. Instead, it studies greenhouse-inspired control benchmarks formulated as QUBO problems and evaluates them with exact, classical, and D-Wave solution workflows Glover et al. (2018); D-Wave Systems Inc. (2026,b). The main benchmark formulation is solved with exact, classical, and D-Wave hybrid methods under a common logical structure, while direct-QPU experiments are performed on reduced instances of the same benchmark class to account for current hardware and embedding limitations D-Wave Systems Inc. (2026b,a). This distinction is important because direct execution on D-Wave quantum hardware requires minor embedding of the logical binary quadratic model into the physical qubit graph, and the quality of this mapping can affect observed solution quality D-Wave Systems Inc. (2026a,b). Accordingly, the contribution of this paper is a practical benchmark of structured greenhouse-inspired QUBO formulations on current D-Wave workflows, with emphasis on feasibility, solution quality, and realistic benchmarking rather than claims of quantum advantage King et al. (2015).

2 Greenhouse-inspired benchmark model

We consider a reduced greenhouse-inspired benchmark with a single binary actuator, namely a heater. At time step k, the control decision is

$$u_k \in \{0, 1\}, \tag{1}$$

where $u_k = 1$ denotes heater on and $u_k = 0$ heater off. The system state consists of the air temperature $T_{\mathrm{air},k}$ and a slower body or canopy temperature $T_{\mathrm{body},k}$.

For the benchmark formulation, the state evolution is described in discrete time by

$$\mathbf{x}_{k+1} = A_d \mathbf{x}_k + B_u u_k + B_{\mathrm{out}} T_{\mathrm{out},k} + B_{\mathrm{sol}} S_k, \qquad \mathbf{x}_k = \begin{bmatrix} T_{\mathrm{air},k} \\ T_{\mathrm{body},k} \end{bmatrix}, \tag{2}$$

where A_d is the discrete-time state-transition matrix, B_u is the heater input vector, B_{out} is the disturbance vector associated with outdoor temperature $T_{\mathrm{out},k}$, and B_{sol} is the disturbance vector associated with solar radiation S_k. The initial condition is denoted by $\mathbf{x}_0$.

Plant performance is represented through a reduced growth proxy. First, an effective growth temperature is defined as

$$T_{\mathrm{grow},k} = \alpha T_{\mathrm{air},k} + (1-\alpha) T_{\mathrm{body},k}, \tag{3}$$

where $\alpha \in [0,1]$ is a weighting factor. Second, a bounded light factor is introduced as

$$L_k = \frac{S_k}{S_k + K_{\mathrm{light}}}, \tag{4}$$

where K_{light} is a light-saturation parameter. The instantaneous growth proxy is then written as

$$g_k = L_k \Big(g_{\max} - \eta \big(T_{\mathrm{grow},k} - T^\star \big)^2 \Big), \tag{5}$$

where $g_{\max}$ is the maximum growth value, η is the curvature parameter, and $T^\star$ is the preferred growth temperature. The growth proxy is not clamped to zero. Therefore, for sufficiently large deviations from $T^\star$, the instantaneous value g_k can become negative. In the reported feasible benchmark solutions, the temperature trajectories remain within the prescribed bounds and do not exploit this negative-growth regime. The proxy should therefore be interpreted as a quadratic performance score rather than a calibrated agronomic growth model.

Energy use and electricity cost are computed diagnostically as

$$E_k = P_{\mathrm{heater}} u_k \Delta t, \qquad C_k = \lambda_E p_k E_k, \tag{6}$$

where P_{heater} is the heater power, Δt is the sampling interval, p_k is the electricity tariff, and λ_E is an energy-cost scaling factor.

Climate feasibility is assessed on the air-temperature trajectory through

$$T_{\min} \le T_{\mathrm{air},k} \le T_{\max} \qquad \forall k, \tag{7}$$

where $T_{\min}$ and $T_{\max}$ denote the admissible lower and upper temperature bounds. The benchmark therefore captures, in reduced form, the trade-off between climate regulation, growth-oriented performance, and energy use, while remaining structured enough for exact, classical, hybrid, and quantum-annealing-based optimization.

Further details of the toy-model construction, including the continuous-time thermal model, exogenous input generation, thermostat baseline, and benchmark parameter values, are given in Appendix A.

3 QUBO formulation

For the optimization problem, the exogenous trajectories $T_{\mathrm{out},k}$, S_k, and p_k are treated as known data, and the heater schedule

$$u_k \in \{0,1\}, \qquad k = 0, \ldots, H-1, \tag{8}$$

is taken as the decision variable. In the benchmark formulation, the no-lag case is assumed, so that

$$u_{\mathrm{applied},k} = u_k, \tag{9}$$

which is consistent with the no-lag case in (A8). This no-lag assumption is used to keep the benchmark focused on the combinatorial scheduling problem rather than actuator-response modelling. In real greenhouse systems, heating dynamics and actuator delays may be relevant; these can be incorporated by replacing (9) with the lagged mapping described in Appendix A.

The optimization problem is written directly in QUBO form as

$$Q(\mathbf{z}) = Q_{\mathrm{growth}} + Q_{\mathrm{energy}} + Q_{\mathrm{switch}} + Q_{\mathrm{bound}}, \tag{10}$$

to be minimized over the binary vector

$$\mathbf{z} = \left(u_0, \ldots, u_{H-1}, \{z^L_{k,m}\}, \{z^U_{k,m}\}\right)^\top \in \{0,1\}^n. \tag{11}$$

Here, Q_{growth} represents the negative growth contribution, so that minimizing it favors temperatures close to the preferred growth regime; Q_{energy} accounts for electricity consumption; Q_{switch} penalizes excessive heater switching; and Q_{bound} enforces the hard air-temperature bounds through quadratic penalties with binary slack variables.

To obtain the individual terms in (10), the discrete-time dynamics (2) are unfolded over the horizon. Since the model is linear and the exogenous inputs are fixed, the state at step k can be written as an affine function of the binary heater sequence,

$$\mathbf{x}_k = \bar{\mathbf{x}}_k + \sum_{t=0}^{H-1} \Gamma_{k,t} u_t, \qquad k = 0, \ldots, H, \tag{12}$$

where $\bar{\mathbf{x}}_k \in \mathbb{R}^2$ collects the contribution of the initial state and the exogenous inputs, and $\Gamma_{k,t} \in \mathbb{R}^{2\times 1}$ is the state sensitivity with respect to the binary control u_t. The air temperature and growth temperature then follow from (2) and (3) as

$$T_{\mathrm{air},k} = \mathbf{e}_{\mathrm{air}}^\top \mathbf{x}_k = \bar{T}^{\mathrm{air}}_k + \sum_{t=0}^{H-1} \gamma^{\mathrm{air}}_{k,t} u_t, \tag{13}$$

$$T_{\text{grow},k} = \mathbf{e}_{\text{grow}}^\top \mathbf{x}_k = \bar{T}_k^{\text{grow}} + \sum_{t=0}^{H-1} \gamma_{k,t}^{\text{grow}} u_t, \tag{14}$$

where $\mathbf{e}_{\text{air}} = [1,\, 0]^\top$, $\mathbf{e}_{\text{grow}} = [\alpha,\, 1-\alpha]^\top$, $\bar{T}_k^{\text{air}} = \mathbf{e}_{\text{air}}^\top \bar{\mathbf{x}}_k$, $\bar{T}_k^{\text{grow}} = \mathbf{e}_{\text{grow}}^\top \bar{\mathbf{x}}_k$, $\gamma_{k,t}^{\text{air}} = \mathbf{e}_{\text{air}}^\top \Gamma_{k,t}$, and $\gamma_{k,t}^{\text{grow}} = \mathbf{e}_{\text{grow}}^\top \Gamma_{k,t}$.

Using (5), (4), and (14), the negative growth contribution in (10) can be written, up to an additive constant, as

$$Q_{\text{growth}} = \eta \Delta t \sum_{k=0}^{H-1} L_k \left(\bar{T}_k^{\text{grow}} + \sum_{t=0}^{H-1} \gamma_{k,t}^{\text{grow}} u_t - T^\star \right)^2. \tag{15}$$

The energy term follows directly from (6) and (9),

$$Q_{\text{energy}} = \lambda_E P_{\text{heater}} \Delta t \sum_{k=0}^{H-1} p_k u_k, \tag{16}$$

and the switching term is written in quadratic binary form as

$$Q_{\text{switch}} = \lambda_S \sum_{k=1}^{H-1} \left(u_k + u_{k-1} - 2u_k u_{k-1} \right). \tag{17}$$

For binary variables, the expression satisfies $u_k + u_{k-1} - 2u_k u_{k-1} = |u_k - u_{k-1}|$, because it is equal to zero when the two consecutive heater states are identical and equal to one when they differ. Thus, (17) penalizes heater switching events while remaining quadratic.

The hard air-temperature bounds in (7) are incorporated through quadratic penalties with binary slack variables. For each $k = 1, \ldots, H$, we impose

$$T_{\text{air},k} - T_{\min} - s_k^L = 0, \qquad T_{\max} - T_{\text{air},k} - s_k^U = 0, \tag{18}$$

with nonnegative slack variables represented in binary form as

$$s_k^L = \sum_{m=0}^{M-1} 2^m \Delta_s z_{k,m}^L, \qquad s_k^U = \sum_{m=0}^{M-1} 2^m \Delta_s z_{k,m}^U, \tag{19}$$

where $z_{k,m}^L, z_{k,m}^U \in \{0, 1\}$, Δ_s is the slack resolution, and M is the number of slack bits.

The two slack variables encode the lower and upper inequalities separately. At any binary assignment where both penalty residuals in (18) vanish, the relation $s_k^L + s_k^U = T_{\max} - T_{\min}$ holds by construction as a consequence of the shared temperature variable. However, this relation is not imposed as an additional constraint. For approximate solvers, one or both residuals may be nonzero simultaneously; in that case, the decoded temperature is penalized but not explicitly corrected by the

slack variables. The formulation should therefore be interpreted as a penalty-based encoding of the two one-sided inequalities. It doubles the slack-bit footprint and may introduce additional coupling structure in the BQM, which is one reason why more compact inequality encodings are identified as future work.

Substituting (13) into (18) yields

$$Q_{\text{bound}} = A_{\text{bound}} \sum_{k=1}^{H} \left(\bar{T}_k^{\text{air}} + \sum_{t=0}^{H-1} \gamma_{k,t}^{\text{air}} u_t - T_{\min} - \sum_{m=0}^{M-1} 2^m \Delta_s z_{k,m}^L \right)^2 + A_{\text{bound}} \sum_{k=1}^{H} \left(T_{\max} - \bar{T}_k^{\text{air}} - \sum_{t=0}^{H-1} \gamma_{k,t}^{\text{air}} u_t - \sum_{m=0}^{M-1} 2^m \Delta_s z_{k,m}^U \right)^2, \tag{20}$$

where $A_{\text{bound}} > 0$ is the penalty weight.

In the numerical benchmarks, we use $M = 6$ slack bits per bound and time step, with slack resolution $\Delta_s = 0.25°\text{C}$. This gives a maximum representable slack value of

$$s_{\max} = (2^M - 1)\Delta_s = 15.75°\text{C}.$$

The bound-penalty weight is set to $A_{\text{bound}} = 120$ for both lower- and upper-bound constraints. These values were chosen to make temperature-bound violations energetically unfavorable while keeping the slack resolution sufficiently fine for the decoded greenhouse trajectories.

We use standard binary slack encoding for transparency and reproducibility. More compact or unbalanced slack encodings may reduce the number of auxiliary variables and improve direct-QPU embedding efficiency; these alternatives are left for future work.

Equivalently, the problem can be written in standard matrix form as

$$\min_{\mathbf{z} \in \{0,1\}^n} \mathbf{z}^\top Q \mathbf{z} + c, \tag{21}$$

where Q is the symmetric QUBO matrix and c is a constant offset. This representation is directly compatible with classical QUBO solvers and quantum annealing hardware Glover et al. (2018); D-Wave Systems Inc. (2026,a).

4 Solvers and benchmark protocol

We consider four main solution workflows for the $H = 24$ benchmark instance, together with additional direct-QPU experiments on reduced instances. As a reference baseline, the decoded control problem is solved exactly by exhaustive enumeration of all binary heater schedules $u \in \{0,1\}^H$. For each candidate schedule, the greenhouse state trajectory is recomputed using the simulator dynamics, and the corresponding objective and feasibility quantities are evaluated directly. This exact reference is computationally restricted to moderate horizons, but it provides the ground-truth optimum for the $H = 24$ benchmark used in this study.

Table 1 Solver and sampler settings used in the benchmark. SA, PIA, hybrid, and direct-QPU workflows are executed using ten independent repetitions. The hybrid settings denote requested Leap Hybrid BQM time limits rather than full end-to-end wall-clock runtimes.

Workflow	Settings
Exact enumeration	Exhaustive enumeration of all 2^H binary heater schedules, followed by simulator-side decoding and objective evaluation.
SA	`SimulatedAnnealingSampler`; 10 seeds; `num_reads`=5000; `num_sweeps`=3000.
PIA	`PathIntegralAnnealingSampler`; 10 seeds; `num_reads`=5000; `beta_schedule_type=custom`; `Hd_field=[10,0]`; `Hp_field=[0,10]`.
Leap Hybrid BQM	`hybrid_binary_quadratic_model_version2p`; 10 independent submissions; requested time limits of 15 s, 30 s, and 60 s. The 30 s setting is used as the representative hybrid result in the main comparison.
Direct QPU	`EmbeddingComposite(DWaveSampler)`; solver `Advantage2_system1`; 10 independent submissions; `num_reads`=5000; annealing time 20 μs; `auto_scale=True`.

The QUBO-based methods solve the binary quadratic model defined in (21). As a classical baseline, we use `SimulatedAnnealingSampler`, which performs thermal annealing on the binary quadratic model. As a simulated quantum annealing baseline, we use `PathIntegralAnnealingSampler`, which implements path-integral simulated quantum annealing on a classical computer.[1] In addition, we use D-Wave's Leap Hybrid BQM workflow, which operates on the same logical BQM while handling decomposition and solver orchestration internally D-Wave Systems Inc. (2026b,b). Direct-QPU experiments are performed separately on reduced instances of the same benchmark class. For direct QPU execution, the logical binary quadratic model must be minor-embedded into the QPU connectivity graph, introducing chains and embedding-dependent overhead D-Wave Systems Inc. (2026a,c,b); King et al. (2015).

The main benchmark instance uses $H = 24$ and $\Delta t = 1$ h, corresponding to one day of hourly binary heater scheduling. The initial conditions are $T_{\mathrm{air},0} = 18.0°\mathrm{C}$ and $T_{\mathrm{body},0} = 17.0°\mathrm{C}$, with a start time of 06:00. With the above slack encoding, this instance yields a BQM with 312 binary variables, consisting of 24 heater variables and 288 slack bits, and 4596 quadratic couplers. For this instance, exact enumeration, classical simulated annealing, path-integral simulated quantum annealing, and the Leap Hybrid BQM workflow are compared under the same decoded objective and feasibility evaluation. The stochastic solvers are executed repeatedly, and the reported values distinguish between best feasible decoded solutions and averages over feasible decoded runs. For the hybrid workflow, a 30 s requested solver time is used as the representative setting in the main comparison. Additional 15 s and 60 s requested time-limit runs are used only as diagnostic checks to test whether the decoded hybrid performance improves monotonically with requested solver time.

[1] We use the term simulated quantum annealing to emphasize that the method is executed classically. Its relationship to quantum-inspired optimization and its comparative advantage over classical simulated annealing are problem-dependent and remain subject to benchmarking.

The D-Wave Ocean components used in the reported runs were:

- `dwave-system` v1.34.0,
- `dimod` v0.12.21,
- `dwave-samplers` v1.7.0.

The `dwave-ocean-sdk` meta-package was not installed as a separate package in the execution environment.

All stochastic workflows are evaluated using ten independent repetitions. This number was chosen as a compromise between repeated-run variability and the practical cost of cloud-based hybrid and QPU access. The resulting statistics are therefore intended as benchmark diagnostics rather than high-confidence estimates of the full solver-output distributions. For this reason, the tables report standard deviations where appropriate, and the text avoids statistical claims about time-limit monotonicity or solver superiority beyond the tested settings.

Direct-QPU experiments are performed for reduced horizons $H = 10$, $H = 12$, and $H = 14$. These instances are small enough for direct embedding and execution on the quantum annealer. The purpose of these reduced direct-QPU experiments is not to replace the main $H = 24$ benchmark, but to assess the behavior of direct quantum annealing on smaller greenhouse-inspired QUBOs under the same decoding and feasibility-evaluation procedure.

For all QUBO-based methods, the optimization variables are the full binary vector in (11), including both the heater variables and the slack bits introduced in (19). Solver comparison, however, is carried out at the level of the decoded heater schedule. After sampling, the control sequence u is extracted from the returned binary vector, the corresponding state trajectory is recomputed using (2), and the decoded solution is re-evaluated using the original simulator quantities defined in (5), (6), and (7). This common post-processing step ensures that all QUBO-based methods are assessed against the same physical model and feasibility criterion rather than only raw BQM energy.

The reported key performance indicators are the decoded objective value, total energy consumption, total growth, air- and body-temperature trajectories, the number of temperature-bound violations, the number of heater switching events, and solver runtime. Total energy consumption is computed as

$$E_{\text{tot}} = \sum_{k=0}^{H-1} E_k, \tag{22}$$

where E_k is defined in (6). Total growth is computed as

$$G_{\text{tot}} = \sum_{k=0}^{H-1} g_k \Delta t, \tag{23}$$

where g_k is defined in (5). The number of temperature-bound violations is defined by

$$N_{\text{viol}} = \sum_{k=0}^{H} \mathbf{1}[T_{\text{air},k} < T_{\min} \ \vee \ T_{\text{air},k} > T_{\max}] \,, \tag{24}$$

and the number of heater switching events is

$$N_{\text{switch}} = \sum_{k=1}^{H-1} |u_k - u_{k-1}|. \tag{25}$$

The decoded simulator-side objective reported in the results is

$$J_{\text{total}}(u) = G_{\text{tot}}(u) - J_E(u) - J_S(u), \tag{26}$$

with

$$J_E(u) = \lambda_E \sum_{k=0}^{H-1} p_k E_k, \qquad J_S(u) = \lambda_S N_{\text{switch}}. \tag{27}$$

This decoded objective differs from the raw QUBO energy $Q(\mathbf{z})$. In particular, the bound-penalty term Q_{bound} is not part of the physical performance objective; instead, temperature-bound satisfaction is evaluated separately through N_{viol}. Therefore, all solvers are compared using the same decoded physical objective and feasibility criterion rather than raw BQM energy.

For instances with a known exact optimum, objective gaps relative to the exact decoded solution are additionally evaluated. For D-Wave hybrid runs, the reported timing refers to the requested hybrid solver time limit and the corresponding D-Wave reported solver time, excluding external queueing, network latency, and notebook overhead. The 15 s, 30 s, and 60 s settings should therefore be interpreted as requested hybrid time limits rather than complete end-to-end wall-clock runtimes. For direct-QPU runs, the reported hardware timing includes the QPU access time returned by the D-Wave system. Because direct-QPU performance depends on embedding and sampling details, these results are reported separately from the main $H = 24$ benchmark.

The purpose of this benchmark is not to establish quantum advantage, but to assess how exact, classical, quantum-inspired, hybrid, and direct quantum-annealing workflows behave on the same structured control benchmark under a common decoding and evaluation procedure.

5 Results

This section reports the benchmark results for the greenhouse-inspired QUBO instances. The main benchmark uses a horizon of $H = 24$ with sampling interval $\Delta t = 1$ h. This instance is solved using exact enumeration, classical simulated annealing (SA), path-integral simulated quantum annealing (PIA), and D-Wave's Leap Hybrid BQM workflow. Direct-QPU experiments are reported separately for reduced

Table 2 Main benchmark comparison for the $H = 24$, $\Delta t = 1$ h instance. The benchmark uses $T_{\mathrm{air},0} = 18.0°\mathrm{C}$, $T_{\mathrm{body},0} = 17.0°\mathrm{C}$, $T_{\min} = 16°\mathrm{C}$, $T_{\max} = 26°\mathrm{C}$, $T^{\star} = 22°\mathrm{C}$, $\lambda_E = 1.5$, and $\lambda_S = 0.5$. For stochastic solvers, Best J, E [kWh], Growth, Switch, and Runtime refer to the best feasible decoded run. The mean objective and standard deviation are computed over feasible stochastic runs. The objective J_{total} is maximized. Exact enumeration is deterministic.

Solver	Feas.	Best J	Mean J $\pm\sigma_J$	E [kWh]	Growth	Switch	Runtime [s]
Exact	N/A	-145.14	-145.14	1100	40.36	11	1028.27
SA	10/10	-150.11	-161.60 ± 6.76	1100	35.39	11	31.84
PIA	10/10	-149.11	-169.94 ± 8.45	1100	35.39	9	3.60
Leap Hybrid BQM, 30 s	7/10	-174.32	-203.84 ± 16.72	1300	41.68	12	29.99

instances with $H = 10$, $H = 12$, and $H = 14$, where direct embedding on the quantum annealer remains practical.

The objective value J_{total} reported in this section is the decoded simulator-side performance measure, where larger values indicate better solutions. Feasibility is assessed through the air-temperature bound constraint, and a decoded solution is counted as feasible only if the number of temperature-bound violations is zero. For stochastic solvers, both the best feasible run and the mean over feasible runs are reported.

Main $H = 24$ benchmark

Table 2 summarizes the main benchmark comparison for the $H = 24$ instance. Exact enumeration provides the reference optimum for this instance. Classical SA and PIA both return feasible decoded solutions in all ten repetitions. The D-Wave Leap Hybrid BQM solver with a 30 s time limit returns feasible decoded solutions in seven out of ten repetitions.

Exact enumeration provides the reference optimum by construction, with $J_{\mathrm{total}} = -145.14$ for the $H = 24$ benchmark. Classical SA reaches a best objective of -150.11, corresponding to an objective gap of approximately 4.97 relative to the exact solution. PIA performs similarly in its best run, reaching -149.11, with an objective gap of approximately 4.03. Both SA and PIA therefore provide feasible near-optimal solutions for this benchmark instance.

The Leap Hybrid BQM result is weaker under the tested configuration. Its best feasible decoded solution reaches $J_{\mathrm{total}} = -174.32$, corresponding to a gap of approximately 29.18 relative to the exact optimum. The mean feasible hybrid objective, -203.84, is also substantially below the corresponding SA and PIA means. Therefore, for this $H = 24$ instance, the tested hybrid workflow does not outperform the classical annealing baselines.

The hybrid result is nevertheless informative. The feasible hybrid runs satisfy the temperature constraints and have zero violations. The performance loss is therefore not primarily a feasibility failure, but rather a weaker trade-off between growth, energy use, and switching. In particular, the best hybrid solution achieves a larger growth proxy than the exact solution, 41.68 compared with 40.36, but does so by using more

Table 3 Diagnostic comparison of Leap Hybrid BQM results for the $H = 24$ instance using different requested hybrid time limits. Only feasible decoded runs are included in the best and mean objective values. The "No feasible" column counts submissions for which no feasible decoded heater schedule was retained under the post-processing criterion.

Hybrid setting	Feasible	No feasible	Best J_{total}	Mean J_{total}	Best E_{tot} [kWh]	Mean E_{tot} [kWh]
15 s, 10 repeats	5/10	5/10	-181.88	-194.35	1300	1360
30 s, 10 repeats	7/10	3/10	-174.32	-203.84	1300	1400
60 s, 10 repeats	2/10	8/10	-185.32	-200.77	1400	1500

energy, 1300 kWh compared with 1100 kWh, and by producing a worse overall objective. This indicates that the decoded hybrid solution remains physically meaningful but is not optimal under the chosen objective weighting.

Since $\Delta t = 1$ h and $P_{\text{heater}} = 100$ kW, the total energy value can be cross-checked directly from the number of heater-on intervals: each heater-on step contributes 100 kWh. Thus, the exact solution energy of 1100 kWh corresponds to 11 heater-on intervals, while the best hybrid solution energy of 1300 kWh corresponds to 13 heater-on intervals.

The penalty weight A_{bound} was selected empirically to prioritize feasibility of the decoded temperature trajectory. This choice makes temperature-bound violations energetically unfavorable, but it can also reshape the BQM energy landscape and affect the relative balance between growth, energy, and switching terms. The weaker hybrid performance should therefore not be interpreted as an intrinsic limitation of hybrid quantum-classical optimization. It may also reflect the present penalty scaling, the dense coupling structure induced by the rollout dynamics and slack variables, and the mismatch between internal BQM energy minimization and the decoded simulator-side objective.

Hybrid time-limit diagnostic

To test whether the weaker hybrid result was caused by insufficient requested solver time, the $H = 24$ hybrid experiment was evaluated using requested time limits of 15 s, 30 s, and 60 s. Table 3 compares the 15 s, 30 s and 60 s hybrid settings.

The hybrid time-limit diagnostic should be interpreted cautiously because each setting was evaluated using only ten independent hybrid submissions. The 30 s setting gives the highest feasible-run rate and the best hybrid objective among the tested settings. However, the non-monotonic behavior from 15 s to 30 s to 60 s should not be interpreted as evidence that longer hybrid runs are generally worse. Rather, it shows that, for this formulation and sample size, increasing the requested hybrid time limit did not produce a reliable improvement in decoded simulator-side performance.

The infeasible hybrid submissions in Table 3 correspond to submissions for which no feasible decoded heater schedule was retained under the post-processing criterion. Therefore, the diagnostic reflects both the internal BQM search and the external decoded-feasibility filter. A larger number of independent hybrid submissions would

be required to distinguish stochastic sampling effects from systematic time-limit behavior.

For this reason, the 30 s setting is retained as the representative hybrid result in the main benchmark comparison, while the 15 s and 60 s settings are treated as diagnostic time-limit checks.

Temperature trajectories

Figure 1 compares the air-temperature trajectories for the best feasible decoded solutions of the benchmark solvers. All reported feasible solutions remain within the admissible temperature interval $[T_{\min}, T_{\max}]$, consistent with the zero-violation counts in Table 2. The exact, SA, and PIA solutions follow broadly similar feasible trajectories, whereas the hybrid solution shows a different heater-timing pattern that leads to higher energy use and a lower objective value.

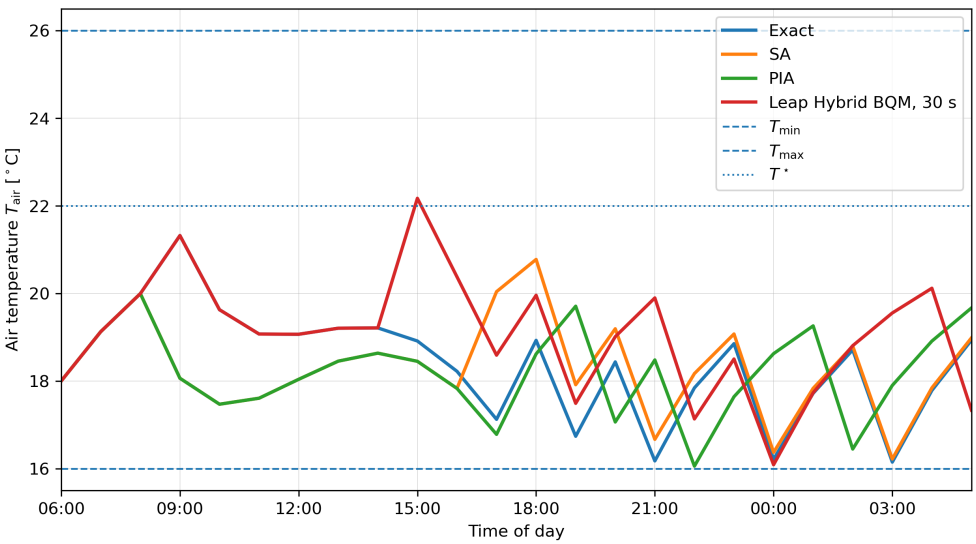

Fig. 1 Best feasible air-temperature trajectories for the $H = 24$ benchmark. Only the best feasible decoded solution of each solver workflow is shown to keep the comparison readable. The admissible interval is defined by $T_{\min} = 16°\text{C}$ and $T_{\max} = 26°\text{C}$.

The corresponding body-temperature trajectories and cumulative-growth curves are provided in Appendix B. The body-temperature trajectories are smoother and less discriminative than the air-temperature trajectories because of the slower thermal response of the second state. Similarly, the cumulative-growth curves mainly support the quantitative growth–energy trade-off already reported in Table 2. They are therefore included as supplementary diagnostics rather than central results.

Table 4 Reduced-instance benchmark results for $H = 10$, $H = 12$, and $H = 14$. The table compares exact enumeration, SA, PIA, and direct QPU execution on the same reduced instances. Stochastic means and standard deviations are computed over feasible decoded runs. All stochastic runs were feasible. QPU exact hits and QPU access time are reported only for direct-QPU rows.

H	Method	Vars.	Coupl.	Best J	Mean $J \pm \sigma_J$	QPU hits	QPU access [s]
10	Exact	130	1005	8.19	–	–	–
10	SA	130	1005	8.19	8.19 ± 0.00	–	–
10	PIA	130	1005	8.19	8.19 ± 0.00	–	–
10	QPU	130	1005	8.19	-1.04 ± 12.71	5/10	0.85
12	Exact	156	1362	-6.61	–	–	–
12	SA	156	1362	-6.61	-6.61 ± 0.00	–	–
12	PIA	156	1362	-6.61	-6.79 ± 0.38	–	–
12	QPU	156	1362	-6.61	-20.53 ± 18.07	2/10	0.88
14	Exact	182	1771	-34.61	–	–	–
14	SA	182	1771	-34.61	-35.19 ± 0.49	–	–
14	PIA	182	1771	-35.11	-36.74 ± 2.33	–	–
14	QPU	182	1771	-35.11	-58.83 ± 23.38	0/10	0.89

Reduced-instance direct-QPU results

Direct-QPU experiments were performed on reduced instances with $H = 10$, $H = 12$, and $H = 14$. These cases are small enough to allow direct execution on the D-Wave quantum annealer after embedding. The purpose of these experiments is not to solve the largest benchmark instance, but to assess how direct QPU execution behaves on reduced greenhouse-inspired QUBOs under the same decoding and feasibility-evaluation procedure.

Table 4 summarizes the direct-QPU results. For each horizon, the exact objective is reported together with the best and mean decoded QPU objective over ten repetitions. The exact-hit count denotes the number of QPU repetitions that recovered the exact decoded optimum.

For $H = 10$, SA, PIA, and the direct QPU all recover the exact optimum in their best runs. However, the mean QPU objective is substantially below the exact value, indicating a broad distribution of decoded QPU solution quality despite the 5/10 exact-hit rate. For $H = 12$, the QPU still recovers the exact optimum in two out of ten repetitions, while SA and PIA remain close to the exact solution in both best and mean performance. For $H = 14$, no QPU exact hit is observed. The best QPU solution has a small objective gap of approximately 0.50, but the mean QPU gap is approximately 24.22, indicating low reliability at this instance size. The contrast between the near-optimal best QPU sample and the much lower mean QPU objective indicates a highly dispersed decoded-solution distribution with high variance, especially at $H = 14$, whereas PIA remains comparatively concentrated around near-optimal decoded objectives.

The reduced-instance results therefore show that direct QPU execution can produce feasible decoded solutions for all tested horizons, but its reliability decreases with horizon and remains weaker than the classical annealing baselines in mean decoded objective. The degradation with increasing horizon appears mainly as a loss of optimality reliability rather than a loss of feasibility.

Discussion of solver behavior

The combined results lead to three observations. First, the greenhouse-inspired QUBO benchmark produces physically interpretable decoded solutions across all solver workflows. Feasibility can be evaluated directly in the simulator, and all feasible decoded solutions satisfy the imposed temperature bounds.

Second, for the exactly benchmarkable $H = 24$ instance, the classical annealing baselines are strong. Both SA and PIA produce feasible near-optimal solutions in all repetitions and outperform the tested Leap Hybrid BQM configuration. The hybrid workflow should therefore not be presented as superior for this benchmark instance. Its value in the present study is instead diagnostic: it shows that D-Wave's scalable hybrid BQM workflow can be applied to the structured greenhouse QUBO, but that its performance is sensitive to the formulation, scaling, and solver interaction.

A limitation of this comparison is that $H = 24$ is still a moderate benchmark size and is not the regime for which large-scale hybrid decomposition workflows are primarily designed. The choice of $H = 24$ was deliberate because it allows comparison against a known exact decoded optimum. However, this also means that the present results should not be interpreted as a definitive assessment of the Leap Hybrid BQM solver on larger multi-day greenhouse horizons. Larger instances, such as $H = 72$ or $H = 168$, should be studied in future work using strong classical heuristics as baselines, since exact enumeration is no longer practical in that regime.

Third, the direct-QPU results provide a complementary hardware-facing benchmark. Direct QPU execution is feasible for reduced instances and can recover exact optima for $H = 10$ and $H = 12$. At $H = 14$, the QPU no longer recovers the exact optimum in the tested repetitions, although the best decoded solution remains close to optimal. This supports a cautious interpretation: for the tested reduced instances, direct QPU execution produced feasible decoded solutions and occasionally recovered the exact optimum, but its mean decoded performance remained weaker than the classical annealing baselines. The tested instances therefore do not support a claim of quantum advantage.

The hybrid time-limit diagnostic is also limited by the number of independent submissions. Each time-limit setting was evaluated with ten repetitions, which is sufficient to reveal non-monotonic behavior in this particular experiment but not sufficient to determine whether the effect is systematic. The 15 s, 30 s, and 60 s comparison should therefore be interpreted as a diagnostic observation rather than a statistically robust statement about hybrid runtime scaling.

Overall, the results support the role of this greenhouse-inspired QUBO as a structured benchmark for annealing-based optimization workflows. They also show the importance of comparing quantum and quantum-hybrid workflows against strong classical baselines under the same decoded objective and feasibility criteria.

6 Limitations

The benchmark is intentionally simplified. It contains a single binary heater actuator and does not model ventilation, humidity, CO_2, crop-specific physiology, actuator saturation, forecast uncertainty, or real greenhouse data. The growth term is a quadratic proxy and should not be interpreted as a calibrated biological growth model.

The main hybrid comparison is restricted to $H = 24$, where exact enumeration is still possible. This choice enables comparison against a known decoded optimum, but it does not probe the large-scale regime for which hybrid decomposition workflows are primarily designed. Larger multi-day horizons should therefore be studied in future work using strong classical heuristics and mathematical-optimization baselines.

The QUBO formulation uses penalty-based bound enforcement with binary slack variables. The penalty value $A_{\text{bound}} = 120$ was selected empirically to prioritize feasible decoded trajectories, but no systematic penalty-sensitivity study is included. Different penalty scalings or alternative inequality encodings may change solver behavior, especially for hybrid and direct-QPU workflows.

The direct-QPU experiments are restricted to reduced horizons $H = 10$, $H = 12$, and $H = 14$ because direct hardware execution requires minor embedding of the logical BQM. These experiments should therefore be interpreted as reduced embedding-aware hardware tests, not as evidence of practical quantum advantage.

The direct-QPU experiments use the default `EmbeddingComposite` behavior and do not manually tune the chain strength. Chain-break fractions, physical-qubit counts after embedding, and embedding-quality statistics were not recorded systematically for the reported runs. Consequently, the present results cannot separate degradation due to increasing logical problem difficulty from degradation due to embedding quality or chain-breaking effects. This is a limitation of the hardware-facing analysis and should be addressed in future direct-QPU benchmarks.

Finally, the benchmark uses a coarse sampling interval of $\Delta t = 1$ h. This is adequate for a structured proof-of-concept benchmark but not for operational greenhouse climate control, where shorter control intervals and actuator-response dynamics would be required.

7 Conclusion

This paper introduced a greenhouse-inspired control benchmark formulated as a quadratic unconstrained binary optimization problem and evaluated it using exact enumeration, classical simulated annealing, path-integral simulated quantum annealing, D-Wave's Leap Hybrid BQM solver, and direct QPU execution on reduced instances. The aim was not to propose a deployable greenhouse controller or to claim quantum advantage, but to construct an application-motivated benchmark that exposes the practical behavior of current annealing-based optimization workflows on a structured control problem.

The results show that the benchmark is useful for separating feasibility, decoded solution quality, and hardware scalability. For the $H = 24$ instance, exact enumeration provides the reference optimum. Classical SA and PIA produce feasible near-optimal solutions in all repetitions, with best objectives close to the exact solution. Under the

tested settings, the Leap Hybrid BQM workflow does not outperform these classical baselines. In the limited ten-repetition diagnostic, increasing the requested hybrid time limit from 30 s to 60 s did not improve the decoded solution quality. This suggests that the limitation is not merely insufficient runtime, but may be related to QUBO scaling, penalty tuning, formulation density, or solver-specific behavior.

The reduced direct-QPU experiments provide a complementary hardware-level result. Direct QPU execution recovers the exact optimum for $H = 10$ and $H = 12$, with exact-hit rates of 5/10 and 2/10, respectively. At $H = 14$, no exact hit is observed under the tested settings, although the best decoded QPU solution remains close to the exact optimum. All decoded QPU solutions are feasible and have zero temperature-bound violations. Thus, the main degradation with increasing horizon is a reduction in optimality reliability rather than a loss of feasibility.

Taken together, the results support a cautious interpretation of current quantum annealing workflows for greenhouse-inspired control QUBOs. Direct QPU execution provides a reduced embedding-aware hardware test, while classical annealing methods remain stronger baselines for the tested instance sizes. Hybrid quantum-classical workflows remain relevant for larger or more complex instances, but the present results show that their performance cannot be assumed and must be benchmarked carefully against classical heuristics.

Future work should investigate improved QUBO scaling, alternative penalty formulations, post-processing of decoded samples, constrained quadratic models, and richer greenhouse models with additional control variables such as ventilation, shading, humidity, and CO_2. The present benchmark uses a coarse sampling interval of $\Delta t = 1$ h, which is sufficient for a structured proof-of-concept benchmark but not intended to represent the full temporal resolution required in operational greenhouse climate control. Practical controllers would require shorter control intervals to capture intra-hour dynamics and actuator constraints.

Appendix A Full greenhouse toy-model specification

This appendix provides the full specification of the greenhouse-inspired toy model underlying the benchmark in Section 2. The benchmark is intentionally simplified, but it retains a two-state thermal structure, weather and solar forcing, a thermostat-based baseline policy, and diagnostic growth and energy calculations.

Continuous-time thermal model

The underlying thermal dynamics are defined in continuous time as

$$\begin{aligned} \frac{dT_{\text{air}}}{dt} &= -(k_{\text{env}} + k_{\text{air,body}})T_{\text{air}} + k_{\text{air,body}}T_{\text{body}} \\ &\quad + k_{\text{env}}T_{\text{out}} + K_{\text{heater}}u_{\text{applied}} + K_{\text{solar}}S, \\ \frac{dT_{\text{body}}}{dt} &= k_{\text{body,air}}T_{\text{air}} - k_{\text{body,air}}T_{\text{body}}, \end{aligned} \tag{A1}$$

where T_{air} is the greenhouse air temperature, T_{body} is a slower body or canopy temperature, T_{out} is the outdoor temperature, S is the solar radiation, and $u_{\text{applied}} \in \{0, 1\}$ is the heater action effectively applied to the system. The parameter k_{env} denotes the heat-exchange coefficient between the greenhouse and the outside environment, $k_{\text{air,body}}$ the coupling coefficient from air to body temperature in the air balance, $k_{\text{body,air}}$ the relaxation coefficient governing how the body temperature follows the air temperature, K_{heater} the heater gain, and K_{solar} the solar-gain coefficient.

The continuous-time model is converted to discrete time by zero-order hold, which yields the discrete-time form used in the main paper, cf. (2). The benchmark is initialized from

$$\mathbf{x}_0 = \begin{bmatrix} T_{\text{air},0} \\ T_{\text{body},0} \end{bmatrix}, \tag{A2}$$

where $T_{\text{air},0}$ and $T_{\text{body},0}$ denote the initial air and body temperatures.

Exogenous input profiles

The exogenous inputs are outdoor temperature, solar radiation, and electricity price. Solar radiation is prescribed by a daytime half-sine profile,

$$S_k = \begin{cases} S_{\max} \sin\left(\pi \frac{h_k - h_{\text{sunrise}}}{h_{\text{sunset}} - h_{\text{sunrise}}}\right), & h_{\text{sunrise}} \le h_k \le h_{\text{sunset}}, \\ 0, & \text{otherwise}, \end{cases} \tag{A3}$$

where $S_{\max}$ is the peak solar radiation, h_k is the hour of day corresponding to step k, h_{sunrise} is the sunrise hour, and h_{sunset} is the sunset hour.

The outdoor temperature profile is generated as a smooth day–night forcing signal. Let

$$h_k = (h_{\text{start}} + k\Delta t) \bmod 24$$

denote the hour of day. In the benchmark, $h_{\text{start}} = 6$, $T_{\text{out,night}} = 2°\text{C}$, $T_{\text{out,day}} = 8°\text{C}$, $t_{\text{out,min}} = 0$, and $t_{\text{out,max}} = 12.5$. The outdoor temperature is defined by

$$T_{\text{out},k} = \begin{cases} T_{\text{out,night}} + \frac{T_{\text{out,day}} - T_{\text{out,night}}}{2} \left[1 - \cos\left(\pi \frac{h_k - t_{\text{out,min}}}{t_{\text{out,max}} - t_{\text{out,min}}}\right)\right], & t_{\text{out,min}} \le h_k \le t_{\text{out,max}}, \\ T_{\text{out,night}} + \frac{T_{\text{out,day}} - T_{\text{out,night}}}{2} \left[1 + \cos(\pi \xi_k)\right], & \text{otherwise}, \end{cases} \tag{A4}$$

where

$$\xi_k = \begin{cases} \dfrac{h_k - t_{\text{out,max}}}{24 - (t_{\text{out,max}} - t_{\text{out,min}})}, & h_k > t_{\text{out,max}}, \\ \dfrac{h_k + 24 - t_{\text{out,max}}}{24 - (t_{\text{out,max}} - t_{\text{out,min}})}, & h_k < t_{\text{out,min}}. \end{cases} \tag{A5}$$

The electricity tariff follows a weekday three-zone structure,

$$p_k = \begin{cases} 0.08, & 0 \le h_k < 8, \\ 0.12, & 8 \le h_k < 10, \\ 0.18, & 10 \le h_k < 14, \\ 0.12, & 14 \le h_k < 18, \\ 0.18, & 18 \le h_k < 22, \\ 0.12, & 22 \le h_k < 24. \end{cases} \tag{A6}$$

Thermostat baseline

To generate a baseline control trajectory for comparison and illustration, the simulator uses a thermostat with hysteresis. Let $T_{\min}$ denote the lower comfort bound and d_{band} the deadband width. The commanded heater action is

$$u_k = \begin{cases} 1, & T_{\text{air},k} < T_{\min}, \\ 0, & T_{\text{air},k} > T_{\min} + d_{\text{band}}, \\ u_{k-1}, & \text{otherwise}, \end{cases} \tag{A7}$$

where the third case preserves the previous heater command inside the hysteresis band. The model may also include an optional lag between the commanded heater action u_k and the applied heater action $u_{\text{applied},k}$; in the no-lag benchmark setting one sets

$$u_{\text{applied},k} = u_k. \tag{A8}$$

This thermostat is used only to generate reference trajectories in the simulator; it is not the optimization method studied in the main benchmark.

Growth and energy diagnostics

The benchmark uses the reduced growth proxy defined in the main text by (3)–(5). For completeness, the interval growth contribution is

$$G_k = g_k \Delta t, \tag{A9}$$

and cumulative growth is obtained by summing G_k over the horizon.

Likewise, the interval energy use and cost are given by (6), and cumulative energy use and cost are obtained by summing the interval quantities over the horizon.

The number of heater switching events is computed as

$$N_{\text{switch}} = \sum_{k=1}^{H-1} |u_k - u_{k-1}|, \tag{A10}$$

and feasibility is assessed with respect to the air-temperature bounds in (7).

Benchmark instance and overview figure

Figure 1 illustrates a representative simulator instance and the main diagnostic outputs of the toy model. This figure is used to visualize the model dynamics and is not the main $H = 24$ solver benchmark instance.

Appendix B Additional trajectory diagnostics for the $H = 24$ benchmark

Figure 2 shows the body-temperature trajectories corresponding to the best feasible decoded solutions of the $H = 24$ benchmark. Compared with the air temperature, the body temperature evolves more smoothly because of the slower thermal dynamics of the second state. The figure is therefore included as an additional diagnostic rather than as a main result.

Figure 3 shows the cumulative growth proxy for the best feasible decoded solutions. The hybrid solution obtains a comparatively high final growth value, but this is achieved with increased energy consumption. This supports the interpretation that the relevant benchmark criterion is the full decoded objective, not growth alone.

Acknowledgements. This project was funded by the German Federal Ministry for Economic Affairs and Energy (BMWE) under grant no. 03EGTTH024 - QUSMOS.

Statements and Declarations

Funding. This work was supported by the German Federal Ministry for Economic Affairs and Energy (BMWK) under grant no. 03EGTTH024 - QUSMOS.

Competing interests. The authors are affiliated with QUSMOS, which is developing quantum-inspired and quantum-ready optimization solutions for industrial applications. Beyond this affiliation, the authors declare no competing interests directly related to the results reported in this manuscript.

Author contributions. H.A. conceived the study, developed the greenhouse-inspired benchmark formulation, implemented the QUBO model and solver workflows, performed the numerical analysis, and wrote the initial manuscript draft. M.B.Z. contributed to the model development, interpretation of results, and manuscript revision. Both authors reviewed and approved the final manuscript.

Data availability. The benchmark parameter files and result data used to generate the reported tables and figures will be provided as supplementary material during review and are available from the corresponding author upon reasonable request after publication.

Code availability. The code used for the simulations, QUBO construction, solver execution, and benchmark analysis will be provided as supplementary material during review and is available from the corresponding author upon reasonable request after publication.

References

Jiang, M., Shan, K., He, C., Li, C.: Efficient combinatorial optimization by quantum-inspired parallel annealing in analogue memristor crossbar. Nature Communications **14**(1), 5659 (2023) https://doi.org/10.1038/s41467-023-41647-2

D-Wave Systems Inc.: QUBOs and Ising Models. https://docs.dwavequantum.com/en/latest/quantum_research/qubo_ising.html. D-Wave Documentation, retrieved March 29, 2026 (2026)

Glover, F., Kochenberger, G., Du, Y.: A tutorial on formulating and using QUBO models. arXiv preprint arXiv:1811.11538 (2018) https://doi.org/10.48550/arXiv.1811.11538

D-Wave Systems Inc.: Reformulating a Problem for the D-Wave System. https://docs.dwavequantum.com/en/latest/quantum_research/reformulating.html. D-Wave Documentation, retrieved March 30, 2026 (2026)

D-Wave Systems Inc.: Leap Service's Hybrid Solvers. https://docs.dwavequantum.com/en/latest/industrial_optimization/leap_hybrid.html (retrieved April 23, 2026) (2026)

Farhi, E., Goldstone, J., Gutmann, S.: A Quantum Approximate Optimization Algorithm (2014). https://doi.org/10.48550/arXiv.1411.4028 . https://arxiv.org/abs/1411.4028

D-Wave Systems Inc.: Minor Embedding. https://docs.dwavequantum.com/en/latest/quantum_research/embedding_intro.html Accessed 2026-03-28

D-Wave Systems Inc.: Minor-Embedding: Best Practices. https://docs.dwavequantum.com/en/latest/quantum_research/embedding_guidance.html Accessed 2026-03-28

King, J., Yarkoni, S., Nevisi, M.M., Hilton, J.P., McGeoch, C.C.: Benchmarking a quantum annealing processor with the time-to-target metric. arXiv preprint arXiv:1508.05087 (2015) https://doi.org/10.48550/arXiv.1508.05087

Bersani, C., Ouammi, A., Sacile, R., Zero, E.: Model predictive control of smart greenhouses as the path towards near zero energy consumption. Energies **13**(14), 3647 (2020) https://doi.org/10.3390/en13143647

Morcego, B., Yin, W., Boersma, S., Henten, E., Puig, V., Sun, C.: Reinforcement learning versus model predictive control on greenhouse climate control. Computers and Electronics in Agriculture **215**, 108372 (2023) https://doi.org/10.1016/j.compag.2023.108372

Chen, S., Liu, A., Tang, F., Hou, P., Lu, Y., Yuan, P.: A review of environmental control strategies and models for modern agricultural greenhouses. Sensors **25**(5), 1388 (2025) https://doi.org/10.3390/s25051388

Ajagekar, A., Lu, M., Huang, Y., You, F.: Energy-efficient ai-based control of semi-closed greenhouse via deep reinforcement learning and robust optimization. Smart Agricultural Technology **5**, 100289 (2023) https://doi.org/10.1016/j.atech.2023.100289

Mansour, M., Sathyanarayanan, K.K., Sauerteig, P., Streif, S.: Adaptive robust greenhouse climate control: Combining deep reinforcement learning and economic optimization. Applications in Engineering Science **25**, 101327 (2025) https://doi.org/10.1016/j.apples.2025.101327

Lu, M., Gao, P., Li, H., Sun, Z., Yang, N., Hu, J.: An optimization approach for environmental control using quantum genetic algorithm and support vector regression. Computers and Electronics in Agriculture **215**, 108432 (2023) https://doi.org/10.1016/j.compag.2023.108432

Xin, P., Li, B., Zhang, H., Hu, J.: Optimization and control of the light environment for greenhouse crop production. Scientific Reports **9**(1), 8650 (2019) https://doi.org/10.1038/s41598-019-44980-z

D-Wave Systems Inc.: Decomposing Large Problems. https://docs.dwavequantum.com/en/latest/quantum_research/decomposing.html (retrieved April 23, 2026) (2026)

D-Wave Systems Inc.: Samplers. https://docs.dwavequantum.com/en/latest/ocean/api_ref_system/samplers.html. D-Wave Documentation, retrieved March 30, 2026 (2026)

D-Wave Systems Inc.: QPU Solver Parameters. https://docs.dwavequantum.com/en/latest/quantum_research/solver_parameters.html. D-Wave Documentation, retrieved March 30, 2026 (2026)

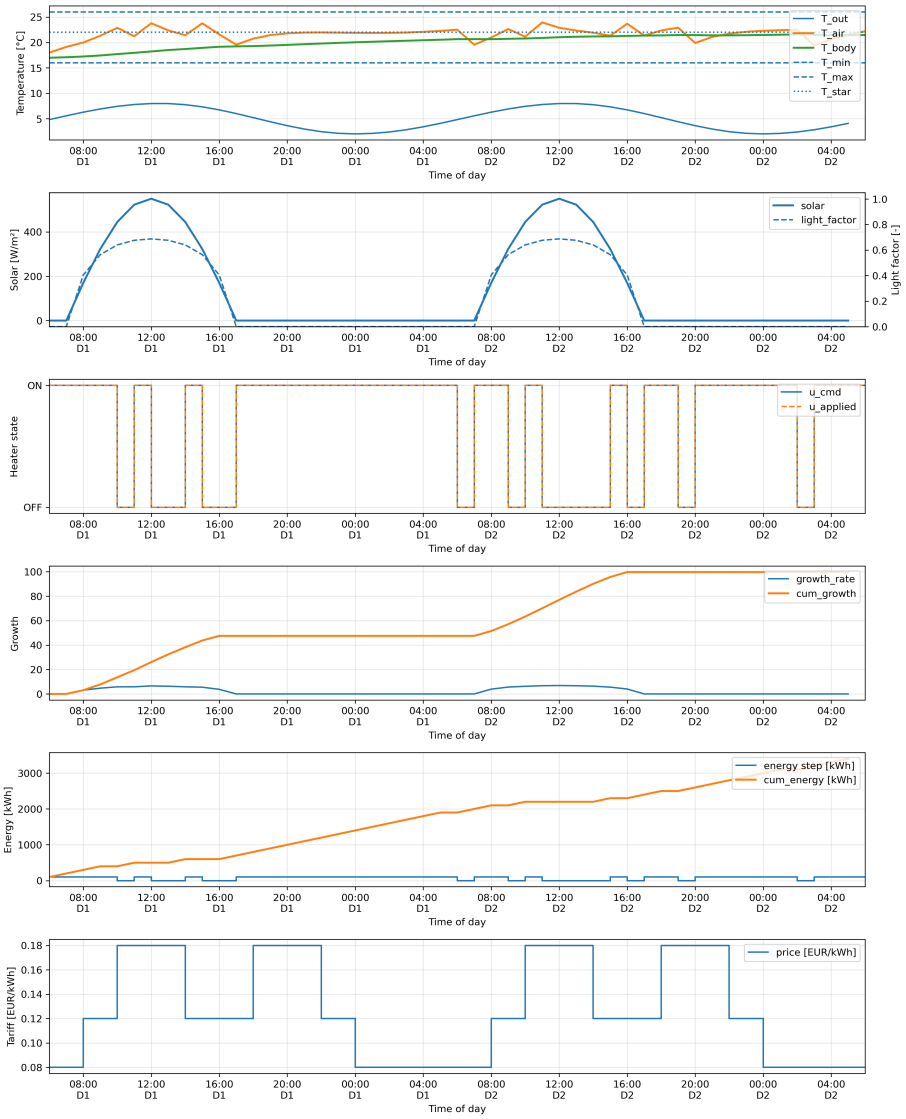

Fig. 1 Overview of the greenhouse toy-model simulator over the benchmark horizon, with the horizontal axis expressed as clock time. From top to bottom, the panels show: (i) the thermal trajectories T_{out}, T_{air}, and T_{body} together with the bounds $T_{\min}$, $T_{\max}$, and the reference temperature $T^{\star}$; (ii) the prescribed solar radiation and the derived light factor; (iii) the heater control variable $u_k \in \{0, 1\}$, where $u_k = 0$ denotes the heater-off state and $u_k = 1$ denotes the heater-on state, shown for both the commanded and applied heater signals; (iv) the instantaneous growth proxy and cumulative growth; (v) the interval energy use and cumulative energy consumption; and (vi) the electricity price tariff used to weight energy consumption over the day. In the plotted simulator overview instance, the horizon is $H = 48$ with $\Delta t = 1$ h, the initial conditions are $T_{\mathrm{air},0} = 18.0°\mathrm{C}$ and $T_{\mathrm{body},0} = 17.0°\mathrm{C}$, the thermal parameters are $k_{\mathrm{env}} = 0.18\ \mathrm{h}^{-1}$, $k_{\mathrm{air,body}} = 0.25\ \mathrm{h}^{-1}$, $k_{\mathrm{body,air}} = 0.06\ \mathrm{h}^{-1}$, $K_{\mathrm{heater}} = 4.0\ °\mathrm{C/h}$, and $K_{\mathrm{solar}} = 0.0045\ °\mathrm{C}/(\mathrm{h\,W\,m^{-2}})$; the thermostat uses $T_{\min} = 16.0°\mathrm{C}$ and $d_{\mathrm{band}} = 1.0°\mathrm{C}$; the solar profile uses $S_{\max} = 550.0\ \mathrm{W\,m^{-2}}$, $h_{\mathrm{sunrise}} = 7.0$, and $h_{\mathrm{sunset}} = 17.0$; the growth model uses $\alpha = 0.4$, $K_{\mathrm{light}} = 250.0\ \mathrm{W\,m^{-2}}$, $g_{\max} = 10.0$, $\eta = 0.18$, and $T^{\star} = 22.0°\mathrm{C}$; and the energy calculation uses $P_{\mathrm{heater}} = 100.0$ kW with $\lambda_E = 1.5$ and a time-dependent tariff p_k.

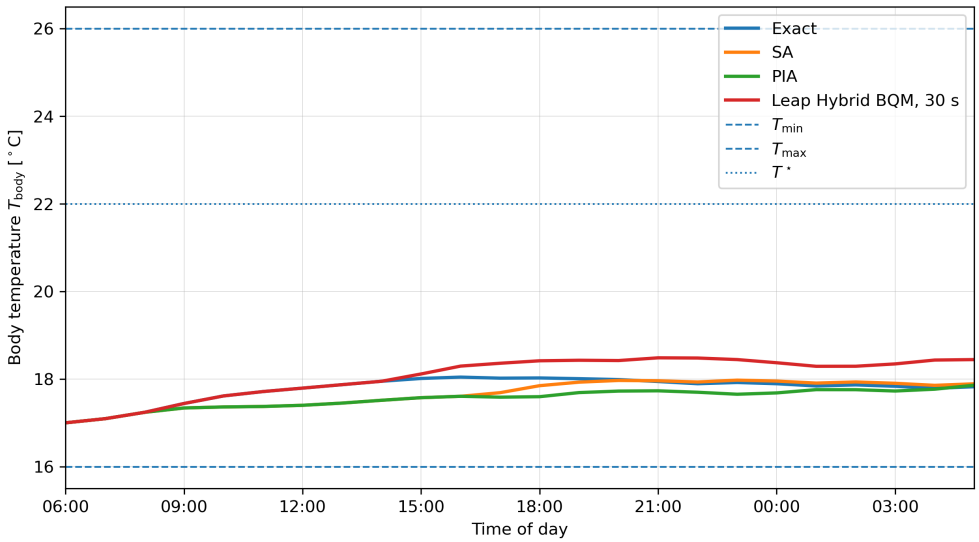

Fig. 2 Body-temperature trajectories for the best feasible decoded solutions of the $H = 24$ benchmark. The body temperature is smoother than the air temperature and mainly serves as a secondary diagnostic of the two-state thermal model.

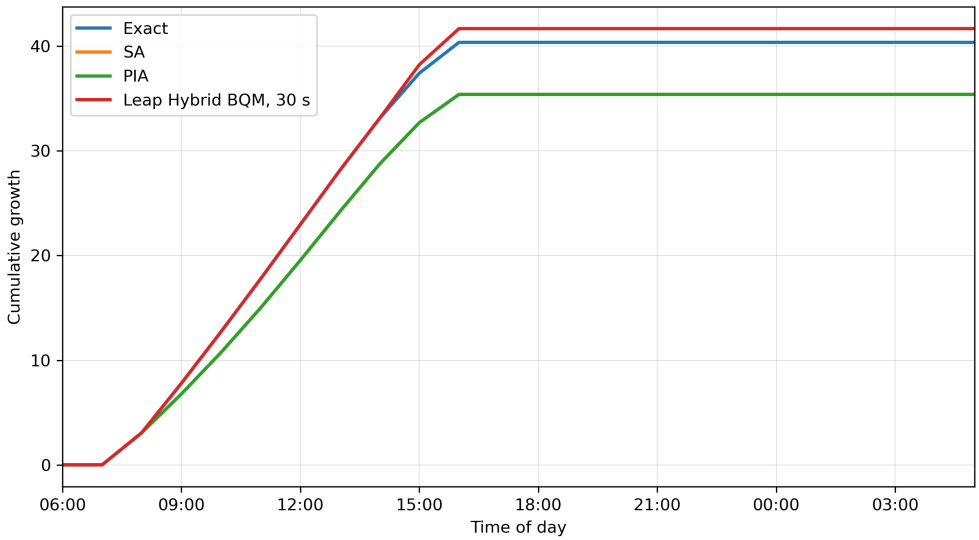

Fig. 3 Cumulative growth proxy for the best feasible decoded solutions of the $H = 24$ benchmark. The figure is provided as an additional diagnostic because the main growth–energy–switching trade-off is reported quantitatively in Table 2.