Cooling consequences of scheduled ventilation under visitor timing uncertainty

LIANG, Xin
Sichuan Bantu Tianxia Caidao Cultural Development Co., Ltd.
liangxin19950218@163.com

Abstract

Comparing a scheduled ventilation strategy with a fixed constant-flow baseline can confuse a scheduling benefit with the choice of capacity. We examine whether midday outdoor-air reductions retain a cooling advantage when constant and timed policies are sized using the same carbon dioxide criterion. A drawing-informed Nanning office prototype was varied across three floor areas and three visitor durations. Four policies were sized with a common periodic mass-balance rule and tested at three visitor arrival times, giving 108 annual EnergyPlus scenarios. In the nominal 18.6725 m2 office with a 60-minute visitor, the constant policy required 12.50 L/s, whereas a two-hour half-flow policy required 19.75 L/s. The latter increased worst-profile annual ideal cooling by 7.37%. Across the nine area-duration cells, the three timed policies required more outdoor-air volume and cooling than their matched constant alternatives. A separately sized, assured-lunch-absence control reversed the result, with the two-hour half-flow policy reducing cooling by 1.05%. Numerical verification retained an unsuccessful coarse-step benchmark: the default 10-minute solver differed from exact transient end states by 28.29 ppm against a preset 25 ppm tolerance. The analytical algorithm matched the independent solution, and fine-step annual diagnostics exposed two short native reports above the research reference. The study provides a bounded comparison of capacity, schedule and visitor assumptions rather than a new control algorithm, a global optimum or a field-validated design. Cooling is thermal demand, not electricity.

Keywords: scheduled ventilation; visitor timing; baseline comparison; carbon dioxide; ideal cooling; numerical verification

1 Introduction

A timer may reduce outdoor air during an assumed lunch absence even when a principal occupant remains or receives a visitor. At a fixed nominal capacity, the resulting reduction in conditioning demand is readily computed. A more useful question is whether that advantage remains when a constant-flow alternative is allowed to change its capacity under the same concentration criterion. A schedule that needs greater peak capacity to tolerate a midday disturbance can deliver more air over the whole workday despite its temporary reduction.

Occupancy- and CO2-based ventilation controls are established. Office-room EnergyPlus comparisons [3], occupancy-based VAV simulations [4], CO2-based minimum airflow control [6], and office experiments [7] already address ventilation and energy trade-offs. Time-averaged VAV control has also been tested in a real building [12]. These systems, baselines and energy metrics differ from the timer-only ideal-load comparison here. Source-strength uncertainty is likewise established through physiological generation studies [2] and sensitivity analysis of ventilation metrics [10]; occupant uncertainty in office energy simulation has been investigated [5].

An initial calculation for this course-design prototype selected a mild timer at 10 L/s that reduced annual ideal cooling by about 0.15%. A subsequent equal-volume constant-flow comparator removed its apparent advantage. That finding motivates the present follow-up. The question is now conditional: when a common sizing rule must cover explicitly specified visitor times, how do required capacity, delivered air volume and ideal cooling change across room sizes and visitor durations? An assured-absence control distinguishes that question from a general claim that timed ventilation is always ineffective.

The contribution is a finite design-response map with explicit comparators, a transparent distinction between assured absence and visitor scenarios, and numerical checks that preserve failed diagnostics. The mass-balance equation and control families are conventional. The study makes no claim to a new prediction algorithm, an empirical probability of failure or a certified ventilation design.

2 Methods

Office prototype and simulation

The room comes from a university design project in Nanning, not a survey of an occupied office. Its nominal net dimensions are 3.85 by 4.85 by 3.8 m, giving 18.6725 m2 and 70.9555 m3. Additional areas of 15 and 30 m2 preserve the plan aspect ratio and height. North and west faces are outdoors; east, south, ceiling and floor are adiabatic approximations. The adjoining open office and its proposed occupants are excluded. The variants are hypothetical geometries, not independent buildings.

Table 1 Central physical and operating assumptions

Input	Value	Evidence
Exterior glazing	WWR 0.8; U 2.7 W/(m2 K); SHGC 0.35	Assumed specification
Opaque enclosure	150 mm concrete; k 1.4 W/(m K); density 2200 kg/m3; cp 900 J/(kg K)	Assumed construction
System and setpoints	Ideal loads; heat 18 C; cool 24 C; dehumidification target 60%	Research assumptions
Availability	Weekdays 09:00–18:00; no holidays	Synthetic operating calendar
Lighting and equipment	5 and 8 W/m2; remain on during lunch	Fixed in all paired policies
People heat and CO2	120 W/person; source 0.0048 L/s/person in primary tests	Upper assumed source, not measured
Outdoor CO2 and leakage	420 ppm; zero base infiltration and door exchange	Simplifying assumptions
Climate	Nanning WMO 594310 CSWD typical-year weather	EnergyPlus weather repository [9]

Honeybee Energy 1.127.2 generated saved HBJSON and IDF definitions; EnergyPlus 26.2.0 performed the annual calculations [1,8]. Heating, cooling and whole-system availability were all restricted to the work period. There was no economizer, heat recovery or demand-controlled ventilation. Ideal supply-air total cooling includes sensible and latent conditioning, but excludes fan electricity, plant efficiency and installation costs. A Sunday-start non-leap synthetic calendar, represented as 2017, contains 260 weekdays; it is not a claim of measured 2017 conditions.

Scenarios and the common sizing rule

The principal occupant remains from 09:00 to 18:00 in the primary cases. One visitor arrives at 11:30, 12:30 or 13:30 and remains for 30, 60 or 120 minutes. Each scenario repeats on every weekday for the annual calculation. The three arrival times are a finite stress set, not a distribution estimated from observations. Every policy is evaluated under the same room, duration and arrival combination.

Policy A supplies a constant rate throughout working hours. B retains 50% during 12:00–14:00; C retains 50% during 12:30–13:30; R90M30 retains 90% during 12:30–13:30. All shut off outside working hours. The design variable is the whole-room outdoor-air capacity at standard density, which is distinct from thermally required supply airflow. Each area-duration cell has one capacity per policy that must cover all three arrival times.

Let x be the excess CO2 concentration above the 420 ppm outdoor background, V the room air volume, N the number of occupants, g the volumetric source per person and q the standard-volume outdoor-air flow. A conventional, well-mixed model [11] gives

$$V\,\frac{dx}{dt} = 10^6 Ng - \alpha q x$$

The sizing model fixes the effective-flow factor alpha at 0.95 for every policy. This is an assumed conservative conversion from standard to effective zone exchange, not a universal value or a parameter fitted to the new runs. For a constant segment, the time constant is volume divided by effective flow; x∞ denotes its equilibrium excess. The exact transition is

$$x(t + \Delta t) = x_\infty + [x(t) - x_\infty]\, e^{-\Delta t/\tau}$$

Exact segment integration gives the 10-minute continuous means. The periodic workday initial concentration is obtained from its affine daily transition, with zero off-hour source and exchange. Bisection over 0.1–30 L/s finds the smallest surrogate capacity whose maximum occupied mean is at most 1000 ppm across all three arrival times. Every capacity is then rounded upward to the next 0.25 L/s. The same search, flow factor, source and rounding apply to constant and timed families. An unrounded-volume check assesses whether the direction is caused solely by capacity rounding.

The 1000 ppm value is a numerical research reference, not a health threshold or a demonstration of regulatory compliance. The study has not established a code-based minimum flow. The sizing criterion is a common surrogate rule; it is not direct global minimization of annual EnergyPlus energy, nor proof that the resulting capacity is the smallest feasible capacity in the full simulator.

Annual comparisons and added checks

The central matrix contains 3 areas by 3 durations by 4 policies by 3 arrivals, or 108 annual runs. EnergyPlus screening uses the maximum occupied concentration reported at 10-minute frequency. Exact continuous means in the sizing model and simulator reports are not assumed to be identical. The data are therefore described as reported concentrations, with reporting and numerical effects examined separately. Both a worst-arrival cooling comparison within a cell and exact-arrival paired differences are retained; no probabilities are assigned to arrivals.

Table 2 Completed calculation stages

Stage	Runs	Purpose
Primary annual matrix	108	Common-rule sizing, three visitor arrival times
Assured-absence annual control	4	Independently size all policies with no visitor and no occupant 12:00–14:00
Joint-input annual probes	6	Two joint assumptions; four runs form exact 12:30-arrival policy pairs
Fine-step annual diagnostics	10	Six planned 2-minute checks; three near-reference checks and one exact-arrival counterpart added
Isothermal short benchmarks	4	Two algorithms by two resolutions against an independent exact solution
Output-only diagnostic repeats	2	Add Detailed output without changing benchmark physics

The assured-absence control uses the nominal area, no visitor and a principal occupant absent from 12:00 to 14:00. All four policies are independently sized by the same rule. Savings are calculated against the constant policy in that same absence scenario; the absolute energy difference between absence and visitor cases is not called a control saving.

Joint assumption J1 combines 0.1 ACH constant background infiltration, source 0.004 L/s/person and glazing U/SHGC of 1.8/0.25. J2 combines zero infiltration, source 0.0048 L/s/person and glazing 3.5/0.50. Capacities remain those of the nominal 60-minute visitor cell. Two policy-specific worst-arrival probes were supplemented with exact-arrival constant counterparts; the reported policy contrasts use

only matching 12:30 arrivals. These are discrete combined assumptions, not exhaustive or probability-based uncertainty analysis. Source variation is independent of fixed 120 W/person heat.

Numerical verification and reporting

Every annual run checks people, actual outdoor-air rate, source generation, zero cooling outside availability, sensible-plus-latent cooling balance and power-time energy consistency. Inputs, weather, engine, logs and losslessly compressed raw outputs are retained. The follow-up protocol was recorded after the original study but before its new primary runs; it is not presented as preregistration. Additional reporting, near-reference and exact-pair diagnostics are explicitly recorded as post-protocol additions. No capacity was retuned after the primary screening.

Six planned 2-minute annual checks cover nominal visitor policies, small-room long visits and the absence control. A coarse dynamic benchmark discrepancy triggered three further C-policy checks, and an additional constant run completed an exact-arrival fine-step pair. Fine reported values are also aggregated into equal 10-minute windows before comparison. Preset common-window tolerances are 20 ppm for the peak and 2% for annual cooling. Native 2-minute values remain a separate diagnostic; they are not silently substituted into the original 10-minute rule.

The independent dynamic benchmark is an all-adiabatic 70.9555 m3 room at 22 C under a clearly labelled artificial constant-weather file. CO2 generation switches between 0, 0.004 and 0.008 L/s, while outdoor air switches between 10 and 5 L/s; the prescribed source adds no people heat. The first saved concentration supplies only an initial condition. All subsequent states are propagated independently with the exact exponential solution and a physical density correction, with no fitted dynamic coefficients. The first day is excluded from error summaries.

The benchmark crosses the default ThirdOrderBackwardDifference and AnalyticalSolution algorithms with 10- and 2-minute steps. Preset end-state tolerances are respectively 25 ppm, 10 ppm and 1 ppm for either analytical run. EnergyPlus registers CO2 as a system-state Average output; that label alone does not establish exact continuous-time integration. Two output-only repeats check Detailed reports against native timestep reports. The benchmark compares both exact end states and exact interval means and preserves an unsuccessful end-state test.

3 Results

What the equal-volume baseline changes

In the initial nominal-source visitor case, the old 10 L/s timer retained 9 L/s for one of nine work hours. Its equal-volume constant comparator is 9.888889 L/s. The timer required 2937.366 kWh thermal and reached 994.133 ppm, whereas the equal-volume constant case required 2936.995 kWh and reached 985.446 ppm. The cooling difference of 0.371 kWh is small; its significance here is that the timed arrangement demonstrated no additional benefit in that comparison. The same direction persisted in the earlier 2-minute pair. Those earlier runs motivate the new matrix and are not counted as new follow-up scenarios.

Capacity and response across visitor scenarios

Figure 1 shows the common-rule capacities. In the nominal 60-minute visitor cell, A, B, C and R90M30 require 12.50, 19.75, 17.00 and 13.25 L/s. The midday reductions must compensate with higher full-flow capacities when a visitor can coincide with their reduction windows. Larger room volume buffers a short source pulse, whereas longer visits require greater exchange. These are transient effects under explicit duration assumptions, not universal area-based minimum rates.

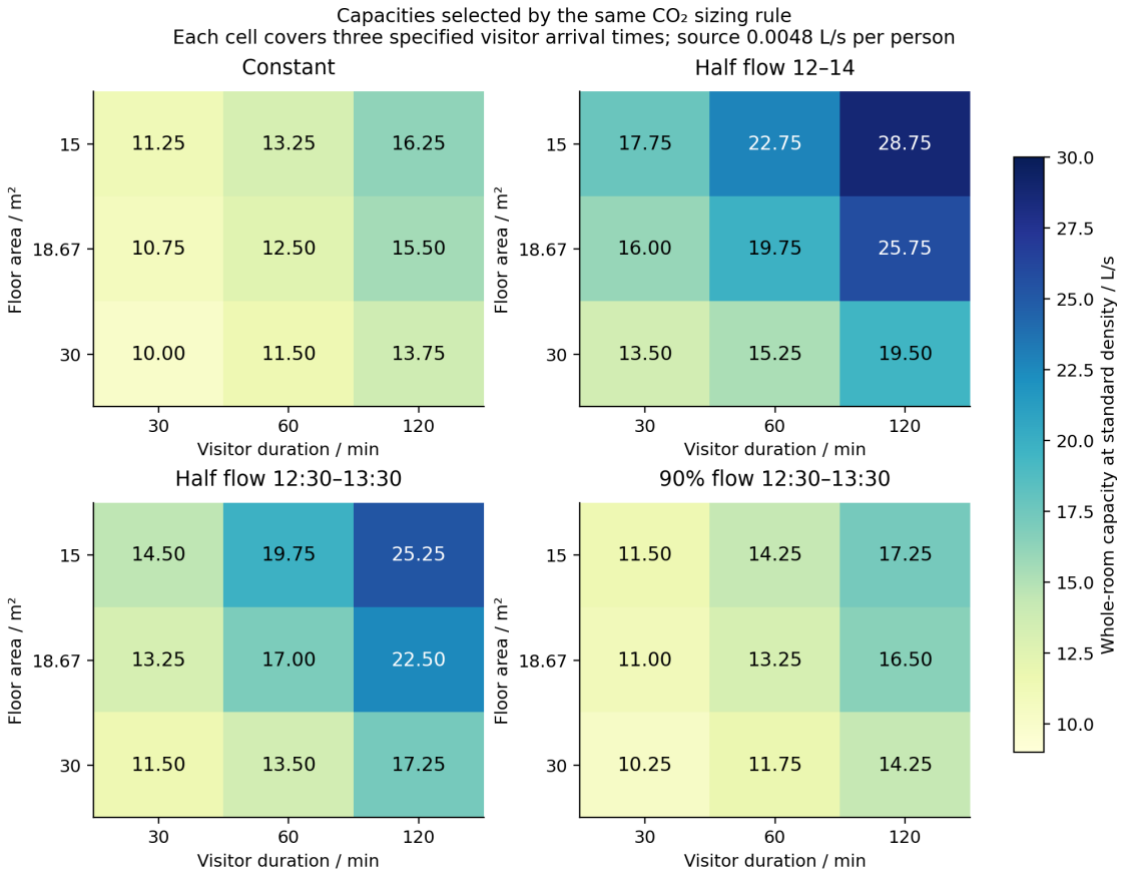

Figure 1 Whole-room standard-density capacities under the common surrogate rule. Each area-duration cell covers three specified arrival times at the upper assumed source. Values are hypothetical research capacities, not certified design rates.

Table 3 Nominal area and 60-minute visitor cell

Policy	Capacity L/s	Worst reported CO2 ppm	Worst cooling kWh thermal	Cooling difference %
A	12.50	982.8	3049.89	+0.00
B	19.75	984.9	3274.56	+7.37
C	17.00	991.1	3206.54	+5.14
R90M30	13.25	979.8	3076.42	+0.87

All 108 primary reported peaks were below 1000 ppm, ranging from 801.2 to 993.6 ppm, so no planned capacity-repair step was used. This is an annual reported-value screen, not proof that every instantaneous concentration remained below the reference. The approximate minimum effective-flow ratio across the annual runs was 0.9807, above the assumed 0.95 sizing factor within this particular model.

Every tested timed policy increased annual ideal cooling relative to its constant alternative under common-rule sizing. Exact-arrival paired increases were 2.38–13.99% for B, 1.02–11.31% for C and 0.14–1.33% for R90M30. Figure 2 shows worst-arrival cooling differences for the nine cells. In the nominal cell these were 7.37%, 5.14% and 0.87%. A reduction during one part of the day did not imply less air or conditioning over the complete workday.

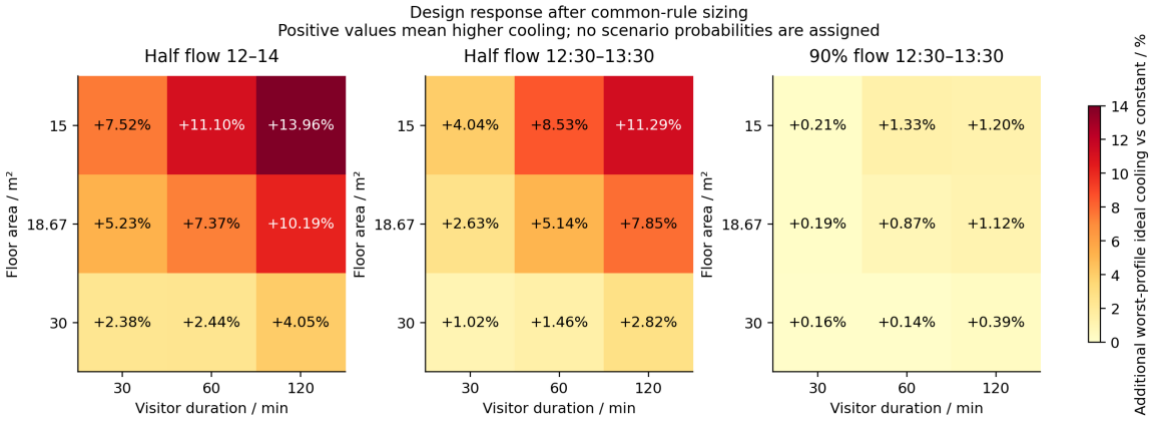

Figure 2 Additional worst-arrival ideal cooling relative to the constant alternative sized by the same rule. Positive values denote a penalty within this finite scenario set. Different cells carry different conditional design capacities.

Daily standard-volume increases were 17.87–57.26%, 8.61–46.75% and 1.04–6.35% for B, C and R90M30. The unrounded surrogate capacities also gave positive timed-volume increases in every cell: 18.15–58.25%, 8.66–48.02% and 1.19–6.05%, respectively. This confirms that the volume direction is not solely a 0.25 L/s rounding effect. It does not establish an unrounded annual-energy optimum. Supplementary Figure S2 shows the nominal visitor-day trajectories.

Assured absence gives a different result

In the independently sized no-visitor absence control, A required 8.50 L/s and 2814.468 kWh thermal. B required 8.75 L/s outside its reduction window but delivered 8.50% less daily outdoor-air volume and used 2784.923 kWh, a 1.05% cooling reduction. C and R90M30 reduced cooling by 0.70% and 0.14%. Figure 3 contrasts these matched effects with the nominal visitor set. The reversal shows the importance of the assumed absence; it does not assign a frequency or expected annual value to either use pattern.

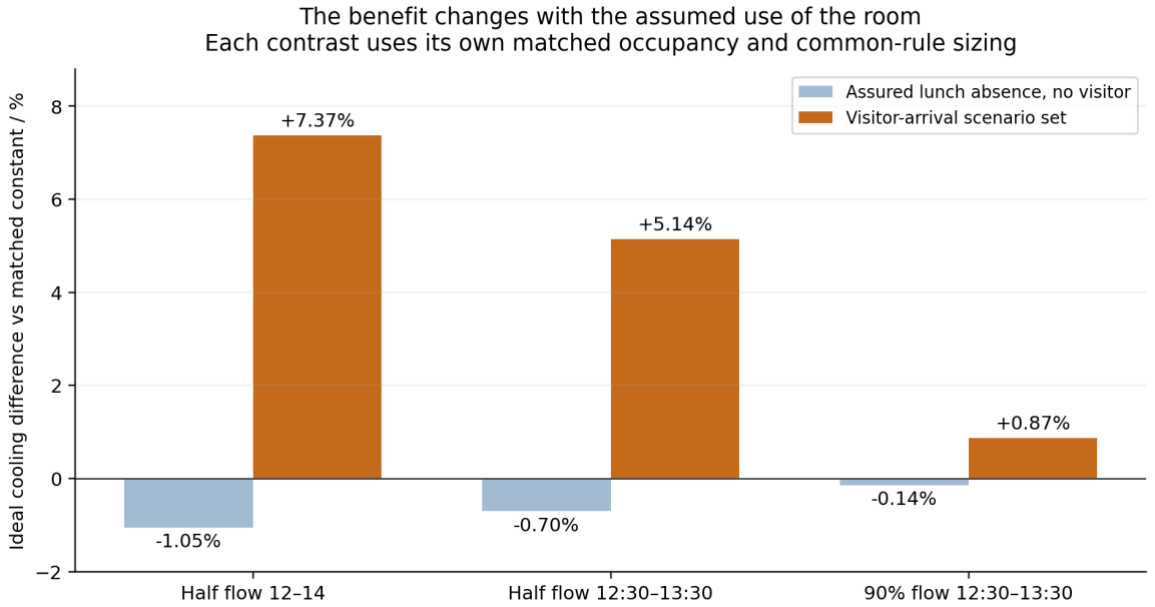

Figure 3 Matched cooling differences under assured absence and the nominal visitor-arrival set. Each contrast uses its own common-rule constant comparator and identical occupancy within that contrast. Absolute loads across the two use patterns are not interpreted as policy savings.

The two exact-arrival joint-input pairs retained the sign for R90M30. Its cooling exceeded A by 0.91% in J1 and 0.74% in J2, although absolute loads changed substantially. Both had reported peaks below the reference. These limited probes support the reported contrast for those assumptions, without establishing robustness to arbitrary simultaneous changes. Supplementary Figure S1 presents their values.

Transient and resolution limits

The analytical benchmark agreed with independently propagated end states to numerical precision. The default algorithm at 10 minutes had a maximum end-state error of 28.29 ppm, exceeding its preset 25 ppm criterion; the failure remains in Table 4. At 2 minutes the corresponding error was 5.94 ppm, within the 10 ppm criterion. The native default reports happened to be closer to the continuous interval means in this switched-source test, with maximum differences of 3.02 and 0.59 ppm. That observation is specific to this benchmark and is not treated as proof that all native reports are exact means.

Table 4 Isothermal benchmark against exact transient end states

Algorithm	Step min	Maximum error ppm	Preset tolerance ppm	Pass
Default third order	10	28.293	25	No
Default third order	2	5.944	10	Yes
Analytical	10	<0.000001	1	Yes
Analytical	2	<0.000001	1	Yes

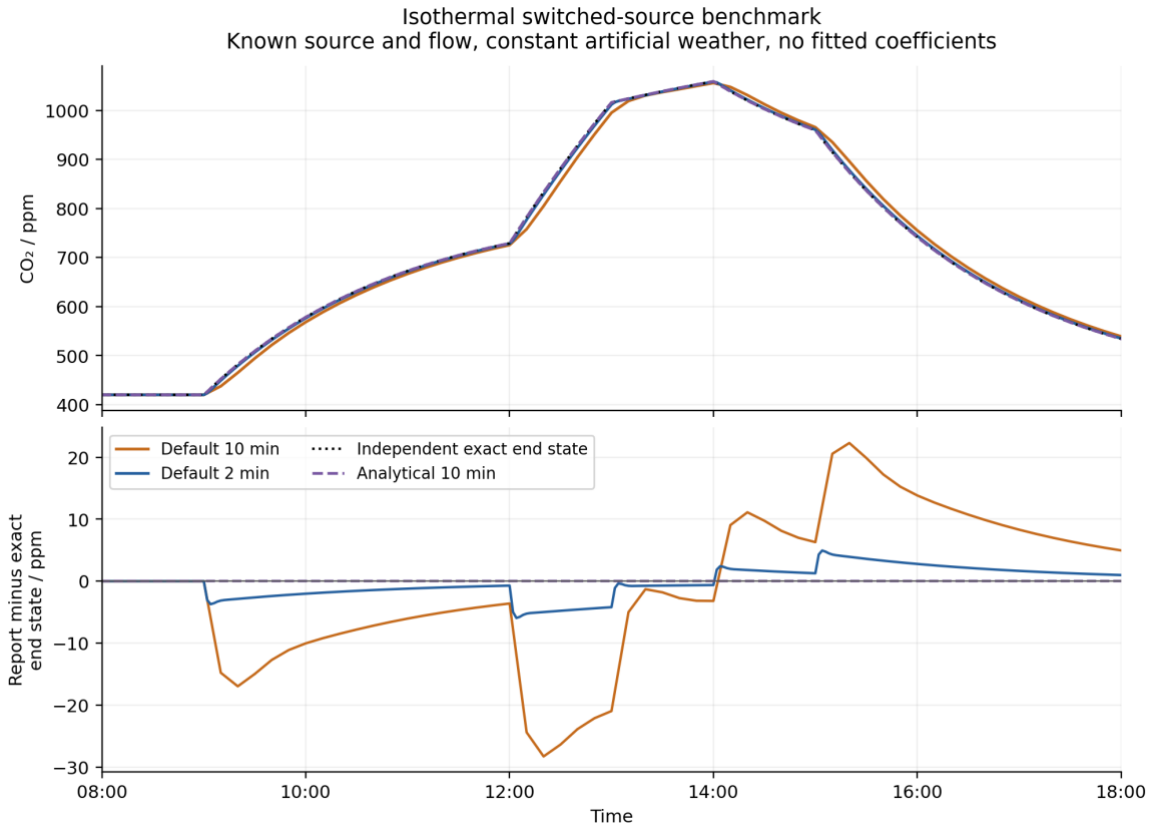

Figure 4 Independent dynamic verification under artificial constant conditions. The coarse default-solver end-state test exceeds its preset tolerance; the figure retains that discrepancy rather than relabelling it as a pass.

The Detailed and native reports had one system report per zone interval in both 10-minute output-only repeats. The observed benchmark differences therefore cannot be explained away as averaging many hidden system reports. They illustrate algorithm and sampling effects around discontinuities. The annual reported-value screen should not be interpreted as an exact continuous-time exposure bound.

All ten annual fine-step checks met the preset common-window tolerances. The largest peak difference after aggregation to 10-minute windows was 11.57 ppm, and the largest annual cooling difference was 0.267%. Exact-arrival nominal B cooling penalties were 7.376% at 10 minutes and 7.307% at 2 minutes; R90M30 penalties were 0.871% and 0.864%. Their signs and approximate magnitudes were stable under this diagnostic.

Two C cases nevertheless exceeded 1000 ppm in their native 2-minute reports: 1002.079 ppm in the 15 m2, 60-minute cell and 1000.121 ppm in the nominal cell. Their corresponding fine outputs aggregated to 10-minute windows peaked at 985.016 and 984.234 ppm. Figure 5 separates the statistics. The difference combines numerical resolution and aggregation; it should not be attributed to averaging alone. Values this close to an assumed line cannot establish practical compliance or safety.

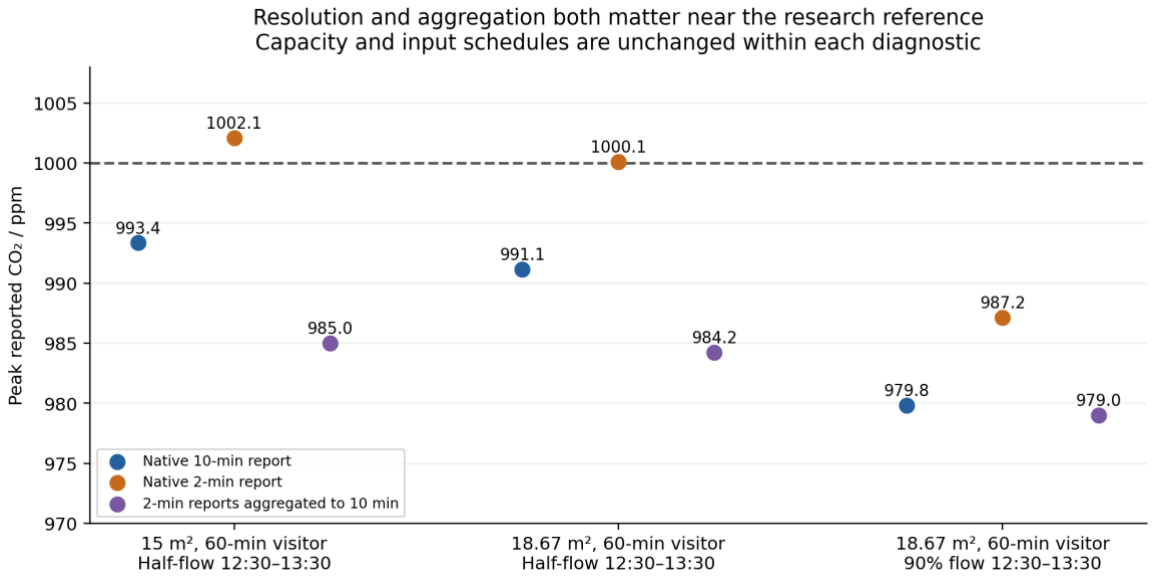

Figure 5 Native coarse reports, native fine reports and equal-window aggregates for near-reference cases. These are distinct reported statistics. The research reference does not certify instantaneous concentration, health or regulatory compliance.

4 Discussion

The main result is conditional on the control family and use envelope. A timer can reduce cooling during an assured absence. When the sizing rule must also cover a visitor arriving during reduced flow, a greater full-flow capacity is required. Delivering that capacity throughout the remaining workday can outweigh the midday reduction. A fair comparison therefore needs both a transparent capacity-setting rule and matching use scenarios; maintaining an arbitrary constant baseline can obscure this mechanism.

The response map is useful because it exposes which assumption drives the comparison. Short pulses in larger volumes are buffered, while smaller volumes and longer pulses require greater exchange. The map does not establish a new physical law: it applies the established transient balance to a finite, reproducible design question. Nor does it establish that constant flow is superior to all possible schedules. A pre-flush, a different capacity-window combination or a feedback controller might change the result and would require its own common-constraint comparison.

The common rule makes all families comparable, but its conservatism and finite grid remain material. Alpha of 0.95 is lower than the simulated exchange ratios, and upward rounding gives unequal residual margins. Capacities minimize a surrogate concentration criterion, not annual energy in the coupled model. The unrounded volume check reduces concern that the observed direction is purely a rounding artefact, while leaving open a finer energy-based capacity search. The largest B/C penalties are more substantial than the small R90M30 differences, which should be interpreted cautiously.

Numerical verification also limits the claim. The exact benchmark verifies an implementation under isolated conditions, not a real building. It revealed a coarse transient discrepancy that successful engine completion alone would miss. Fine-step annual tests supported the cooling contrast but also exposed short reports above the nominal reference. Reporting interval, solver algorithm and aggregation must be stated together; an output labelled Average is not sufficient evidence of exact physical integration. These checks support a screening study with explicit limits rather than a ventilation guarantee.

Transfer to operation is constrained by assumed glazing and construction, zero base leakage, excluded door exchange and the absence of measured calibration. The source range is hypothetical, and source

strength is varied independently of people heat. CO2 is an occupant-related tracer, not a complete measure of indoor air quality. The schedules repeat synthetic events on every weekday and imply no occurrence probability. Geometry variants change facade area and load totals as well as air volume. A single typical-year climate and an ideal-load system cannot establish actual electricity or whole-system economics. Field-informed inputs or a documented external experimental case would be needed before making operational recommendations.

5 Conclusions

For the specified Nanning cellular-office prototype, midday reductions did not retain a cooling advantage when constant and timed families were sized using the same finite visitor scenario rule. Timed standard-air volumes and ideal cooling increased across all nine tested area-duration cells; in the nominal cell the two-hour half-flow policy incurred a 7.37% worst-arrival cooling penalty. In contrast, an independently sized assured-absence control reduced cooling by 1.05%. These results identify the use assumption and baseline capacity as central to the comparison.

Independent transient calculations verified the analytical algorithm and retained a failed coarse default-solver end-state test. Fine annual diagnostics supported the direction of the matched cooling effects, while two native 2-minute C reports crossed the assumed reference. The practical output is a reproducible conditional design-response map and a clearer comparison procedure. It is not a claim of a new control algorithm, global energy optimality, field calibration or certified minimum ventilation.

Data and reproducibility

The study archive contains the follow-up protocol and documented additions, 128 annual simulation records, four short benchmark cases, two additional reporting runs, the common-rule capacity calculations, paired comparisons, numerical diagnostics, figure scripts and the earlier equal-volume comparison. The archived inputs include the Nanning weather file for the annual simulations and a separate synthetic weather file for the constant-condition benchmark. Simulation records include the model inputs, run logs and losslessly compressed raw outputs.

Use of generative AI

OpenAI Codex assisted with drafting and revising the working manuscript, developing simulation and analysis scripts, and preparing figures and document layout. The annual numerical results come from EnergyPlus simulations. Independent mass-balance calculations and selected finer time-step runs were used to check the numerical results as described in Sections 2 and 3.

References

[1] Crawley DB, Lawrie LK, Winkelmann FC, Buhl WF, et al. EnergyPlus: creating a new-generation building energy simulation program. Energy and Buildings. 2001;33:319-331. https://doi.org/10.1016/S0378-7788(00)00114-6

[2] Persily A, de Jonge L. Carbon dioxide generation rates for building occupants. Indoor Air. 2017;27:868-879. https://doi.org/10.1111/ina.12383

[3] Chenari B, Lamas FB, Gaspar AR, da Silva MG. Simulation of Occupancy and CO2-based Demand-controlled Mechanical Ventilation Strategies in an Office Room Using EnergyPlus. Energy Procedia. 2017;113:51-57. https://doi.org/10.1016/j.egypro.2017.04.013

[4] Wang W, Zhang J, Brambley MR, Futrell B. Performance Simulation and Analysis of Occupancy-Based Control for Office Buildings with Variable-Air-Volume Systems. Energies. 2020;13:3756. https://doi.org/10.3390/en13153756

[5] Kang E, Lee H, Yoon J, Cho H, et al. Investigating the influence of uncertainty on office building energy simulation through occupant-centric control and thermal comfort integration. Energy and Buildings. 2024;322:114741. https://doi.org/10.1016/j.enbuild.2024.114741

[6] Lee JM, Lee KH, Moon JW, Lee SH, et al. Developing a Control Strategy for Minimum Airflow Setting Considering CO2 Level and Energy Consumption in a Variable Air Volume System. International Journal of Thermophysics. 2024;45:173. https://doi.org/10.1007/s10765-024-03474-5

[7] Chenari B, Saadatian S, Gameiro da Silva M. Experimental Assessment of Demand-Controlled Ventilation Strategies for Energy Efficiency and Indoor Air Quality in Office Spaces. Air. 2025;3:17. https://doi.org/10.3390/air3020017

[8] EnergyPlus. Version 26.2 documentation, ZoneHVAC IdealLoadsAirSystem and simulation outputs. Accessed 9 November 2026. https://energyplus.readthedocs.io/en/v26.2.0/guides/input-output-reference/1.31-group-zone-forced-air-units.html

[9] EnergyPlus Weather. China Standard Weather Data, Nanning WMO 594310, EPW file. Accessed 9 November 2026. https://energyplus-weather.s3.amazonaws.com/asia_wmo_region_2/CHN/CHN_Guangxi.Zhuang.Nanning.594310_CSWD/CHN_Guangxi.Zhuang.Nanning.594310_CSWD.epw

[10] Oke O, Persily AK. Sensitivity Analysis of CO2 Concentrations as Ventilation Metrics. 43rd Air Infiltration and Ventilation Centre Conference, Copenhagen, 2023. https://www.nist.gov/publications/sensitivity-analysis-co2-concentrations-ventilation-metrics

[11] EnergyPlus. Version 26.2 Engineering Reference, Carbon Dioxide Predictor-Corrector. Accessed 9 November 2026. https://energyplus.readthedocs.io/en/v26.2.0/guides/engineering-reference/2.5-carbon-dioxide-predictor-corrector.html

[12] Kaam S, Raftery P, Cheng H, Paliaga G. Time-averaged ventilation for optimized control of variable-air-volume systems. Energy and Buildings. 2017;139:465–475. https://doi.org/10.1016/j.enbuild.2016.11.059

Supplementary figures

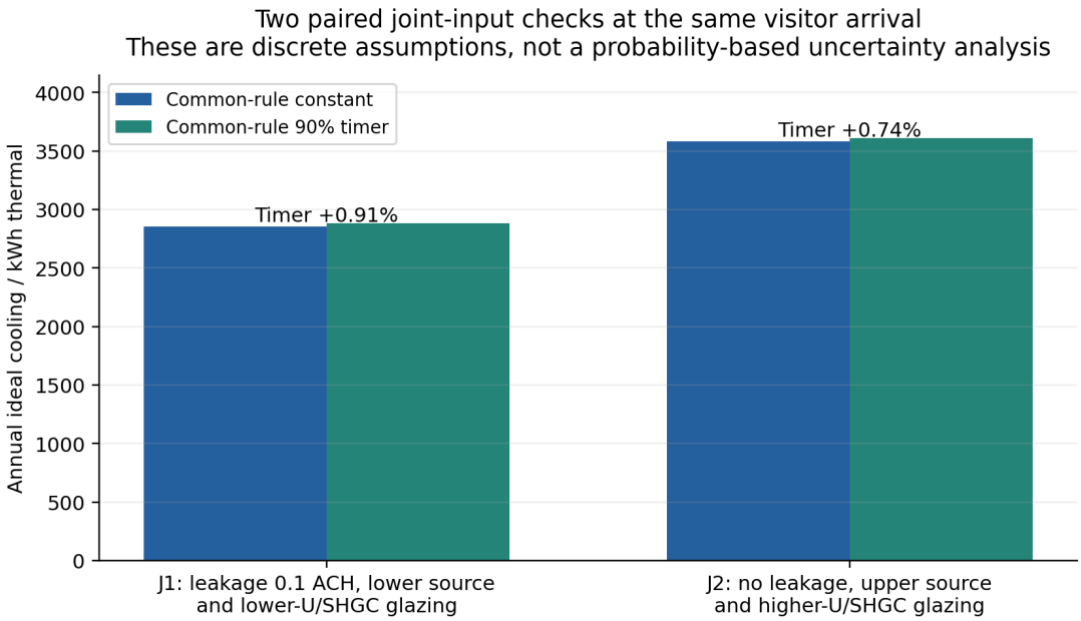

Figure S1 Annual ideal cooling under the two joint input assumptions. Each comparison uses the same 12:30 visitor arrival.

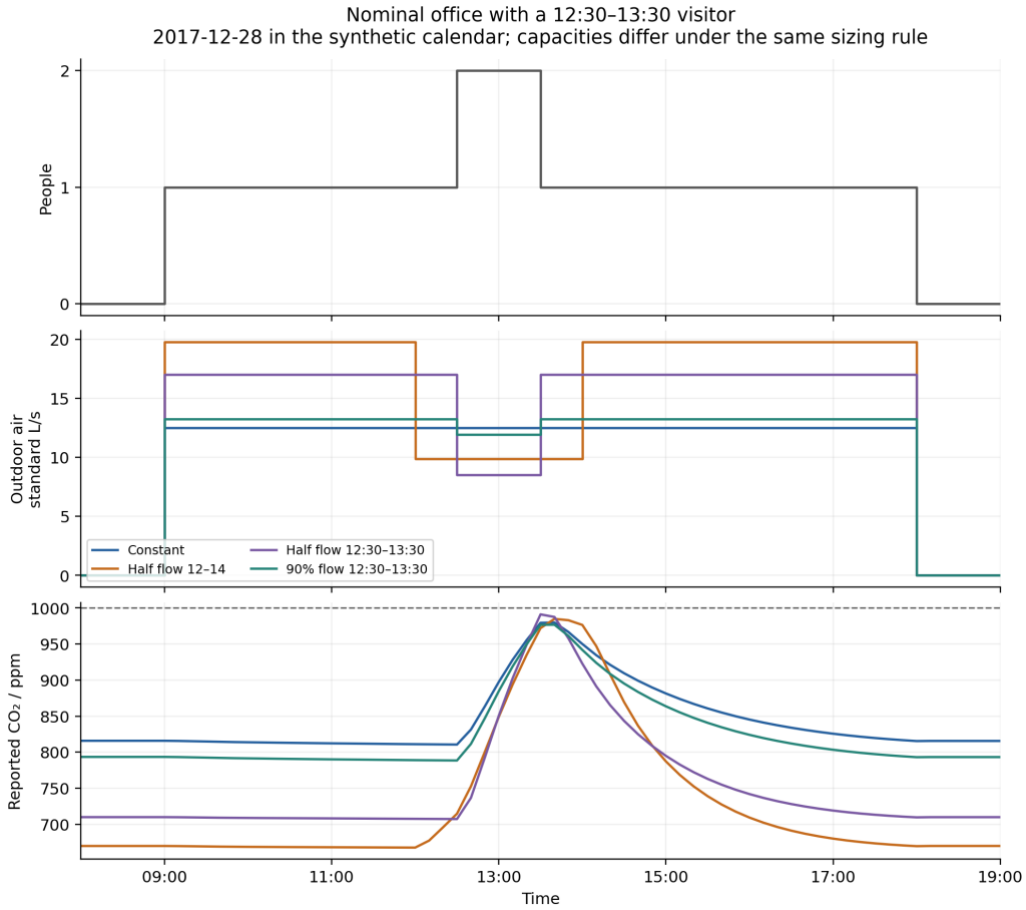

Figure S2 Occupancy, outdoor airflow and reported CO_2 on an example visitor day in the nominal office. The visitor is present from 12:30 to 13:30.