Energy and Buildings 75 (2014) 272–280

Contents lists available at ScienceDirect

Energy and Buildings

journal homepage: www.elsevier.com/locate/enbuild

Using multiobjective optimizations to discover dynamic building ventilation strategies that can improve indoor air quality and reduce energy use

Adams Rackes, Michael S. Waring*

Department of Civil, Architectural and Environmental Engineering, Drexel University, 3141 Chestnut Street, Curtis 251, Philadelphia, PA 19104, United States

ARTICLE INFO

Article history:
Received 5 August 2013
Received in revised form 22 October 2013
Accepted 5 February 2014

Keywords:
Commercial buildings
HVAC systems
Optimal control
Building simulation
Pareto efficiency

ABSTRACT

Ventilation plays a crucial role in promoting the comfort and health of building occupants. It is sometimes costly in terms of energy consumption, but can also be beneficial from an energy perspective when free cooling is available. This work is an exploratory analysis of the hypothesis that simultaneously optimizing energy and indoor air quality (IAQ) objectives can yield better results than existing ventilation control methods. A multiobjective optimization framework was set up to determine optimal time-resolved outdoor airflow and zone temperature setpoints. Test cases were implemented in a modeled office building by numerical, simulation-based optimization for a core and perimeter zone and for typical weather days in January, July, and October in Philadelphia. Results showed that conventional approaches were dominated by the optimized strategies in some cases. Most strikingly, in the core zone in January, mechanical system energy use was reduced by 20–30% with nearly unchanged or improved IAQ. The optimized strategies employed a low-temperature morning flush, a time-shift of some ventilation to the mid afternoon when outdoor air did not require as much heating, and a reduction in ventilation in the evening when it was not as effective at reducing exposure. Cases in July and October demonstrated another benefit: significant IAQ improvements at low energy cost. The results show that there is significant room for improvement in reducing the energy use associated with providing good IAQ.

1. Introduction

Ventilation introduces outdoor air (OA) into a building to dilute indoor-emitted contaminants and protect indoor air quality (IAQ). Most commercial buildings use mechanical ventilation, with rates set in advance according to a standard. For example, the ventilation rate procedure (VRP) of ASHRAE 62.1 specifies a minimum rate V_{bz} for a zone with floor area A_z (m^2) and P_z design occupants as:

$$V_{bz} = R_a A_z + R_p P_z \quad (1)$$

This additive approach includes a per-area component R_a (L/s m^2) intended to dilute non-occupant emissions and a per-occupant component R_p (L/s occ) to provide additional dilution of human bioeffluents [1].

This ventilation air must be conditioned to maintain appropriate ranges of indoor temperatures and relative humidity, which requires energy use by the heating, ventilation, and air conditioning (HVAC) system of the building. HVAC end uses are a major part of energy consumption in commercial buildings, which totaled 18 quadrillion Btu in 2011 alone, accounting for 19% of total US energy use [2]. Some existing strategies modify the rates described by Eq. (1) to achieve energy savings. These include demand controlled ventilation (DCV), which resets the rate based on the actual number of people present when there are fewer than the P_z design occupants, and air-side economizer control, which introduces more OA when it can provide free cooling. These two strategies are not mutually exclusive, and methods have been proposed to combine them more effectively, for example by better identifying and switching between the minimum OA and economizing operating regimes [3].

There is no guarantee, however, that the current standard framework can provide the best IAQ for the least amount of energy. It relies on two local controllers with uncoordinated objectives: one sets a minimum OA rate based entirely on a ventilation standard in conjunction with the design or current zone population, and another determines if a larger rate would be beneficial purely from an energy standpoint. Furthermore, what actually matters for occupant health and productivity is the *concentration* of contaminants

* Corresponding author. Tel.: +1 215 895 1502; fax: +1 215 895 1363.
E-mail addresses: aer37@drexel.edu (A. Rackes), msw59@drexel.edu (M.S. Waring).

http://dx.doi.org/10.1016/j.enbuild.2014.02.024
0378-7788/

indoors over the entire occupied period, not the *rate* at which OA is introduced at a given moment. An alternative approach is developed here: using indoor pollution concentration metrics over an entire day to assess IAQ and the optimization of a combined energy and IAQ objective to determine the best ventilation strategy. Herein, a *ventilation strategy* comprises a trajectory of coupled OA flow rates and zone temperature setpoints. This approach has the potential to identify more flexible control that can take advantage of building and weather dynamics. Although applicable to any building, the information gained from the optimizations is particularly useful for existing buildings, where other interventions like pollutant source control are often not cost-effective.

No one seems to have proposed a similar method for discovering a whole-day, optimized IAQ and energy operational strategy. In the realm of ventilation, the most similar work has focused on optimization in the design phase, including setting static flowrates based on some formulation of optimal tradeoffs between energy and IAQ (or a proxy), e.g., Ref. [4]. There has been less attention to optimal *control* of ventilation over a time horizon. Sherman and Walker have examined dynamic strategies for whole house ventilation in the residential sector, based on a generic contaminant exposure and information about the operation of local exhaust fans [5]. Their work is based on achieving equivalency to a standard, rather than optimization of an outcome, and, moreover, the features of commercial and residential ventilation are quite different.

There has been substantial work on other problems in the optimal control of commercial buildings and their HVAC systems. Two recent reviews provide good summaries [6–8]. Related application areas include single-timestep controller-coordination problems [9], precooling and nighttime ventilation control [10], and active and passive thermal storage [11]. Like this work, most previous studies were conducted in simulated environments. Most were also based on optimizing an objective that was a weighted combination of individual objectives like energy use and comfort [12]; at least one included an instantaneous IAQ constraint or objective [13]. Many optimization routines have been tried and investigated, including standard nonlinear methods like sequential quadratic programming [12] and quasi-Newton algorithms [11], various direct search methods, dynamic programming [14], reinforcement learning [15], and stochastic techniques like genetic algorithm [13,16] and particle swarm optimization (PSO) [17]. For many applications, initial work involving these computationally demanding optimizations eventually led to the development of near-optimal control strategies or simplified rules that could be implemented [18–20]. We hope that solutions to the optimal ventilation strategy problem will follow a similar course.

The present work represents a first step, which is to assess the possible benefits of optimized HVAC control strategies that take into account both energy and IAQ goals. To this end, we formed an IAQ objective whose role is simply to scale concentration metrics so they are comparable to energy costs. We then applied the multi-objective optimization approach to a simulated case study: a small office building in Philadelphia. Two zones were considered independently, on representative days in January, July, and October. We analyzed the transient optimized strategies to identify useful trends in different seasons and spaces. We also used the most efficient tradeoff curves between IAQ and energy use to assess the outcomes of conventional ventilation strategies.

2. Methods

2.1. Optimization problem

The formulation of the finite horizon optimal control problem is well known. There are four classes of variables: system states (x), control variables (u), exogenous variables or disturbances (w), and observed outputs (y). The problem has three fundamental elements: a model that describes the propagation of the states over time and how they are observed, constraints that are imposed by physics or users, and an objective or cost function. (The term "cost" is always used in the context of optimization; it does not refer to monetary value.) In the discrete time formulation, the objective is a function of the outputs, states, control, and exogenous variables at all timesteps 0, . . ., N in the planning horizon:

$$J(y_0, \ldots, y_N,\ x_0, \ldots, x_N,\ u_0, \ldots, u_{N-1},\ w_0, \ldots, w_{N-1}) = J(Y_N, X_N, U_N, W_N) = J(U_N, W_N, x_0) \quad (2)$$

The optimization problem is to find the control trajectory $U_N = \{u_0, u_1, \ldots, u_{N-1}\}$ that minimizes Eq. (2) (in deterministic problems) or its expected value (in stochastic problems), subject to the dynamics imposed by the model and to any additional constraints. The second equality in Eq. (2) is true because the trajectories X_N and Y_N are fully specified by the initial state x_0 and the state model of the dynamics and observations.

To optimize with respect to two goals, the objective can be formed as a linear combination:

$$J(U_N, W_N, x_0) = J_1(U_N, W_N, x_0) + cJ_2(U_N, W_N, x_0) \quad (3)$$

If c is fixed, both objectives can be mapped to a one-dimensional cost space and there is a unique solution. Such an ideal problem would result if one had reliable information about all costs and benefits of ventilation, including health and productivity outcomes of IAQ. However, when costs cannot be established with certainty, as is the case currently with IAQ-related objectives, c parameterizes the family of optimal values known as the Pareto frontier or curve.

2.2. Energy and IAQ cost functions

In this study, the first term in Eq. (3) was selected to be a measure of the HVAC site energy cost:

$$J_1(U_N, W_N, x_0) = 365\frac{\sum_{k=1}^{N} E_{k,\text{fan}}}{A} + 365\frac{\sum_{k=1}^{N} E_{k,\text{cool}}}{A} + 365\frac{\sum_{k=1}^{N} E_{k,\text{heat}}}{A} \quad (4)$$

where the three numerators are, respectively, the total energy consumed on site during the modeled day by the fan, the cooling coil, and the heating coil. The denominator is the floor area served by the system. The objective was formulated with the standard annual energy use intensity (EUI) metric in mind to make it interpretable. The sum can be thought of as the extrapolated energy use intensity (EEUI): the amount of energy per floor area that would be used if every day of the year were exactly the same as the one in the optimization problem.

For this work, we developed a simple IAQ cost function to demonstrate the approach. It should be stressed that more work needs to be done to determine the best form of this objective based on current information about health, productivity, or comfort endpoints. Herein, the IAQ cost is best regarded as a computational tool utilized in the optimization. It was based on concentration metrics for two contaminants: carbon dioxide (CO_2) and total volatile organic compounds (TVOC), which is the sum of individual volatile organics. As a very general rule, CO_2 is an indicator of air quality related to contaminants emitted by occupants, and TVOC is an indicator of air quality related to contaminants emitted by building materials and furnishings. Using contaminants in these two categories is consistent with the current Standard 62.1-2010 VRP philosophy, wherein the per-person ventilation component R_p is required to address occupant-generated contaminants like CO_2

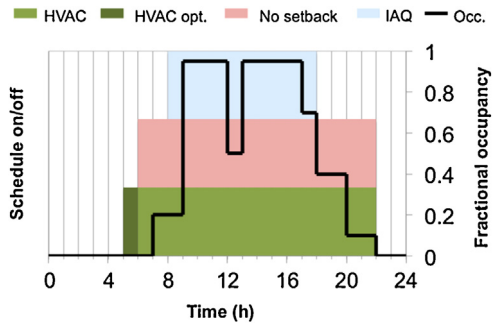

Fig. 1. Testbed schedule values. *HVAC*, *HVAC opt*, *No setback*, and *Occ* are described in Sections 2.3 and 2.4. *IAQ* is the period when contaminant concentrations contributed to the mean and peak IAQ metrics.

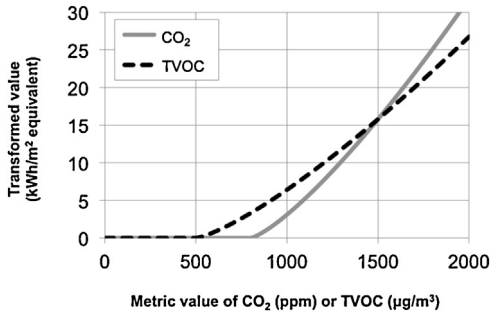

Fig. 2. Transformations for CO_2 and TVOC metric values to implement thresholds and scale to equivalent energy units.

and the per-area component R_a is required to address building-emitted contaminants like VOCs. Neither indicator is perfect by any means, but together they are likely to give a good qualitative picture of IAQ due to indoor-generated pollutants in most mechanically ventilated commercial buildings.

Two metrics were included for each contaminant, the daytime mean and peak concentrations, where "daytime" was defined as 8:00 a.m.–6:00 p.m. This period is depicted as "IAQ" in Fig. 1. There were thus four IAQ-related quantities: mean CO_2, peak CO_2, mean TVOC, and peak TVOC. Each was processed with a transformation of the form:

$$\zeta_{mtr}^{(j)} = \gamma^{(j)} \left(max \left(C_{mtr}^{(j)}, C_{thr}^{(j)} \right) - C_{thr}^{(j)} \right)^n \quad (5)$$

where $\zeta_{mtr}^{(j)}$ is the transformed value of the metric *mtr* for the contaminant *j*, $C_{mtr}^{(j)}$ is the untransformed concentration-based metric, $C_{thr}^{(j)}$ is a compound-specific lower threshold value below which no penalty is imposed, $\gamma^{(j)}$ is a compound-specific scaling constant, and $n > 1$ is an exponent.

The most important parameters in Eq. (5) are the threshold values. The CO_2 threshold was 800 ppm. This relatively low threshold was selected in part because a recent study demonstrated significantly worse occupant mental function and performance at 1000 ppm CO_2 than at 600 ppm, suggesting that CO_2 may be a primary pollutant [21]. For comparison, 1100 ppm is a typical DCV maximum CO_2 concentration setpoint, and levels of up to about 2500 ppm result from following the ventilation rates in ASHRAE 62.1-2010 for some spaces with high occupant densities [1]. The TVOC threshold was 500 $\mu g/m^3$. There are no official standards or guidelines for chronic low-level TVOC exposure in the United States. The German Federal Environment Agency recommended exposure tiers in 2007; the selected threshold is half the 1000 $\mu g/m^3$ point that divides their Level 2 (no "relevant health-related concerns") and their Level 3 ("some objections and distinct health issues") [22]. Given the thresholds, the scaling constants were defined so that a metric value of 1500 $\mu g/m^3$ for TVOC or 1500 ppm for CO_2 would correspond to a transformed value of 15.8 kW h/m^2 (5.0 k Btu/ft^2) equivalent. The exponent *n* was chosen to preserve differentiability but still yield a reasonably linear transformation. Fig. 2 depicts the resulting IAQ cost function transformations.

Finally, the four transformed metrics were combined into a single objective with a weighted sum:

$$J_2(U_N, W_N, x_0) = \sum_{j,mtr} \gamma_{mtr} \zeta_{mtr}^{(j)} \quad (6)$$

We selected weights of 0.25 for the peak metrics and 0.75 for the mean metrics. The intent was to place primary emphasis on exposure, which scales with the mean concentration, while imposing some penalty on solutions with high peaks. The parameters selected for use in Eqs. (5) and (6) are summarized in Table 1.

Table 1
Summary of parameters to transform metrics into equivalent kW h/m^2 and combine them into a scalar cost.

mtr, *j*	$C_{thr}^{(j)}$	*n*	$\gamma^{(j)}$	γ_{mtr}
Mean CO_2	800 ppm	1.3	0.00316	0.75
Peak CO_2	800 ppm	1.3	0.00316	0.25
Mean TVOC	500 $\mu g/m^3$	1.3	0.00199	0.75
Peak TVOC	500 $\mu g/m^3$	1.3	0.00199	0.25

2.3. Building model

The building dynamical model was implemented in EnergyPlus, a validated, physics-based platform with a large support community, predefined models, and detailed weather data [23]. The simulation testbed was a 511 m^2 post-1980 vintage office building in Philadelphia, with four perimeter zones and one core zone each served by a packaged constant air volume (CAV) air handling unit (AHU) with a gas-fired heating coil and a direct expansion cooling coil. The building was based on a US Department of Energy Commercial Reference Building Model [24,25]; a detailed building description with assumptions and data sources can be found in the references.

The model was modified in the following ways for this project:

The fractional occupancy schedule was altered slightly (truncated in the very early morning and late evening) so that the building was not occupied during any periods when the HVAC system was off. It is depicted as "Occ" in Fig. 1.

The unpressurized or nighttime infiltration rates were reduced. Constant values of 0.35 h^{-1} in the perimeter zones and 0.1 h^{-1} in the core zone were used, yielding a building total infiltration rate of about 0.28 h^{-1}, which is consistent with past estimates and field measurements [26]. The infiltration rate during the day was half the nighttime unpressurized value (not one quarter as in the Reference Building model).

Contaminant processes were modeled, according to the differential equation for a well-mixed zone of air: $dC/dt = S - (V_{bz}/V + \lambda_{inf})(C - C_{out})$, where *S* is the contaminant emission source, *V* is the volume of the zone (m^3), λ_{inf} is the infiltration air exchange rate (h^{-1}), and C_{out} is the outdoor concentration. For CO_2, *S* varied with time because it was emitted by occupants. The CO_2 emission rate per occupant was 0.31 L/min, a typical value for office workers, and the outdoor CO_2 concentration was steady at 400 ppm. For TVOC, *S* was constant, representing emissions from building materials and furnishings and neglecting VOC emissions from occupants or secondary reactions with ozone on building surfaces. It was 350 $\mu g/m^3$ h on a volume-normalized basis, which was the

75th percentile in a recent statistical analysis of emissions across the office sector [27] where TVOC represented the sum of 43 individual VOCs. The outdoor TVOC concentration was 69 $\mu g/m^3$, the median from the same analysis. For TVOC, this model also neglected losses due to gas-phase reactions as well as any sorption effects.

The simulation timestep was reduced from 10 to 5 min for greater accuracy.

2.4. Optimization approach

The optimization engine was GenOpt, developed at Lawrence Berkeley National Laboratory, and the optimization algorithm was the included implementation of Hooke and Jeeves' generalized pattern search (GPS) [28]. Such a direct search method is appropriate for simulation-based optimization where objective function derivatives are unavailable or excessively noisy [29]. Furthermore, in this implementation, the convergence tolerance of the simulation engine is taken into account by means of an adaptive precision parameter [30].

To test the supervisory approach, batch optimizations were conducted on three representative days: a cold winter heating day (January 10), a hot summer cooling day (July 26), and a mild fall day near the balance point (October 26), all selected for representativeness from the typical meteorological year (TMY3) data for the NE Philadelphia airport. The time horizon was 1 day, implemented cyclically (the simulations were actually run for 3 days with identical weather, and only the outputs from the third day were used) to prevent the possibility of discovering a control trajectory that performed better on 1 day at the expense of yielding worse results the next. Also, optimizations were separately conducted for the core zone and one of the perimeter zones (the south zone), and only controls for the zone under investigation were altered. Each day/zone combination is called a case.

For each case, the exogenous variable sequence W_N in Eq. (2) represented fixed parameters, and the objective was a deterministic function of the control variables. These were the zone temperature setpoints and OA ventilation rates between 6 a.m. and 10 p.m. In Fig. 1, this time period is the combined area of "HVAC" and "HVAC opt." The control timestep was 15 min, a length used in a similar project [16], leading to a total of 128 variables in the optimization. One limitation was that the system fan always had to operate starting at 6 a.m., even if a later starting time might have performed better. (Attempts to use a hybrid PSO and GPS algorithm with binary fan on/off variables proved intractable.) A crucial step in getting the optimizations to converge was selecting good initial guesses. An approach that proved successful was using the actual zone temperatures and ventilation rates obtained from an initial simulation with standard economizer and DCV control, augmented by a fixed morning flush of ventilation air prior to occupancy.

The problem constraints were implemented primarily through hard box constraints on the input variables. The OA fractions, which define the tradeoff between the ventilation rate and recirculation air rate and thus the amount of ventilation air delivered, had to be between 0 and 1. The zone temperature setpoints had to be between 15 and 28 °C from 6 a.m. to 7 a.m. They had to be between 21 and 24 °C from 7 a.m. to 10 p.m. ("No setback" in Fig. 1). There was also a steep penalty added to the objective for any setpoint-not-met hours to assure that the zone temperature actually achieved the setpoint. During the times not considered in the optimization, from 10 p.m. to 6 a.m., the temperature was allowed to float in the deadband defined by the nighttime setback setpoints, and the OA fraction was zero.

In order to establish the Pareto curve, optimizations were conducted for five values of c in the multiobjective cost of Eq. (3): 0, 0.5, 1, 2, and 5. In addition to these five optimized strategies for each case, two baseline control strategies were defined for comparison. The first (B1) was a standard strategy that used fixed minimum ventilation rates calculated according to the ASHRAE 62.1-2010 VRP. The second (B2) represented an "advanced" commercial OA strategy combining a differential enthalpy economizer and occupancy-tracking DCV, using ASHRAE 62-2001 rates. In both of these baselines, the HVAC system operated from 7 a.m. to 10 p.m. ("HVAC" in Fig. 1). All baselines had the same setback periods and temperature constraints as the optimizations.

3. Results and discussion

3.1. Transient characteristics

The optimized strategies differed from the conventional ones in the most useful and informative ways in January and October. Figs. 3 and 4 show the ventilation, TVOC concentration, coil energy, and in one case CO_2 concentration for selected cases of most interest. Each plot includes three optimized strategies and the two baselines, and where present, vertical lines indicate the IAQ metric period. For ventilation, air changes per hour (ACH, h^{-1}) are shown to facilitate normalized comparisons. We do not include plots of CO_2, fan energy, or zone temperature setpoints because they are not noteworthy. Optimized strategies' CO_2 trajectories were not as varied as the TVOC ones, and were generally similar in form to those resulting from B2. Fan energy was essentially constant with this CAV system and fixed schedule (save the additional first hour of operation in all optimized strategies). The zone temperature setpoints were in fact important for reducing energy use, but changes in them were subtle and often only for a small portion of the day.

The *January core zone* optimized strategies were characterized by three broad trends (Fig. 3). First was a pre-occupancy ventilation flush in the morning. The flush drove down elevated overnight TVOC concentrations (Fig. 3b), reducing both peak and mean TVOC levels during the day. Remarkably, the flush resulted in very little additional heating because it was coupled with a low zone temperature setpoint, about 15 °C. This effect can be observed by the low coil gas use between 6 a.m. and 7 a.m. in Fig. 3c.

The second trend was that ventilation was reduced after the flush, as the zone required sensible heating to be brought up from its nighttime setback to its daytime occupied temperature. (The zone temperature remained at the minimum, 21 °C, for the rest of the day.) One consequence was reduced heating loads in the mid-morning, clearly visible between 8 a.m. and noon in Fig. 3c. Over the morning, ventilation increased, as OA became less costly or even unnecessary to heat, due to warmer outdoor temperatures and the buildup of lagged internal heat gains. This time-shifting of ventilation to later in the day was, as might be expected, greatest with $c = 1$. Under that strategy, it resulted in noon CO_2 levels of around 1000 ppm, or about 100 ppm higher than under either baseline strategy. With $c = 5$, there was less mid-morning reduction, which kept the CO_2 concentration nearly identical to the baselines. For $c = 5$, there was also a second morning OA flush from about 7:30 a.m. to 8:30 a.m. Taken together, the optimized results suggest that in the winter there is significant value to some ventilation time-shifting, and also to staggering the periods of greatest ventilation and greatest sensible heating.

The third trend was that the optimized strategies reduced ventilation late in the afternoon and the evening. Indeed, no optimized strategy, including $c = 5$, found ventilating the building at high rates in the evening to be a useful operation. Ventilation at that time of day simply was not very effective at reducing mean and peak metrics, and ventilating less at that time saved significant heating energy. Only B1, with a fixed ventilation rate, required heating in the evening, which means that that entire heating load

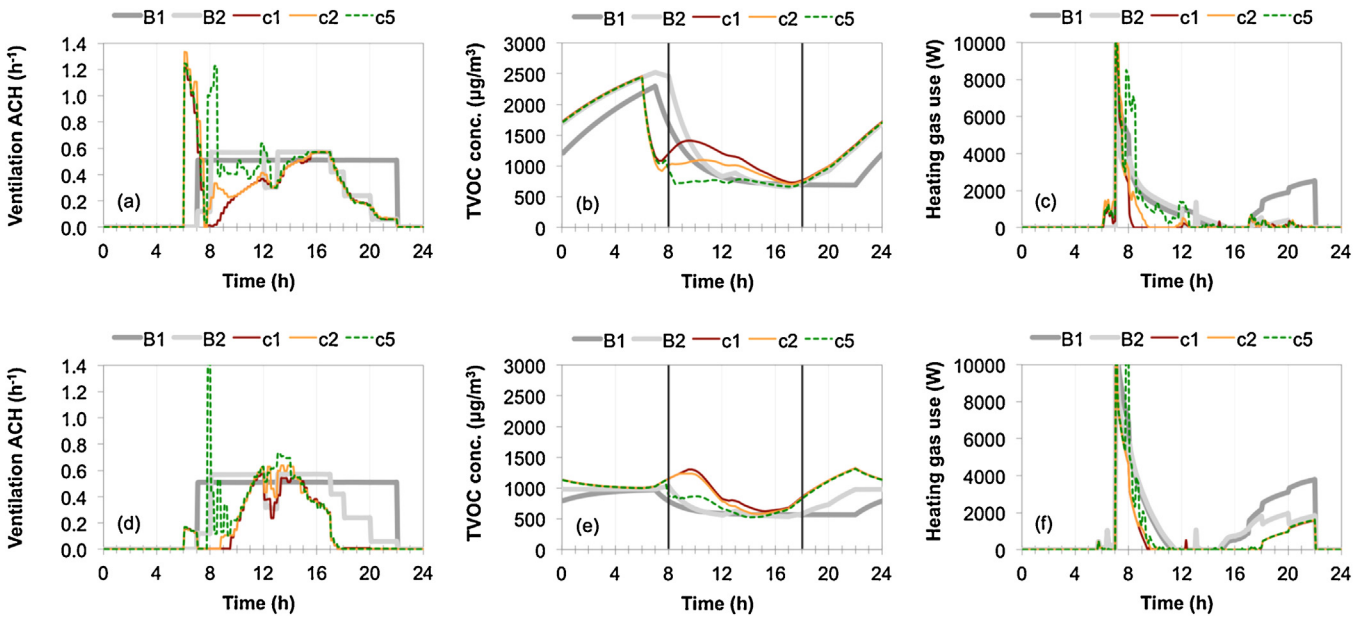

Fig. 3. January results in the core zone (top row) and perimeter zone (bottom row). The columns are: ventilation air changes per hour (ACH), TVOC concentration, and heating gas consumption.

is attributable to ventilation. It is worth noting that, of the three trends described in January in the core zone, the third was quite similar to occupant-tracking DCV, but the first two were very different from both baseline strategies.

January perimeter zone results differed from those in the core primarily because of the much more meaningful infiltration in the perimeter. Infiltration reduced the primacy of ventilation as a dilution mechanism and also limited the nighttime buildup of VOCs (Fig. 3e). Because of this, the results in the perimeter exhibited even more aggressive time-shifting and an earlier reduction in ventilation in the afternoon. These changes, respectively, greatly reduced the morning and evening heating needs (Fig. 3g). The differences between the perimeter and the core are very clear from a comparison of Fig. 3a and d. Unlike in the core, the afternoon OA curtailment occurred much earlier in the day than under occupant-tracking DCV (B2). Nor did the optimized strategies in the perimeter include very significant pre-occupancy morning flushes. Only the strategy with $c = 5$ had anything like one, and this occurred most vigorously during 15 min before the IAQ metric period began at 8 a.m. Essentially, OA was only provided in large amounts in the middle of the day when it was thermally useful to prevent overheating in this south-facing zone.

The *July core zone* optimized strategies were less novel, and did not improve much on the economizer-with-DCV baseline B2. The reason was that, though it was a hot day, it was cool enough to economize with an integrated economizer until about 8:15 a.m. This allowed B2 to save energy, and, as a positive but unintentional consequence, drive down TVOC concentrations. In both the core and perimeter zones, the optimized strategies did bring in OA earlier, at 6 a.m., which reduced the zone temperature to the maximum daytime setpoint of 24 °C an hour earlier than the baselines did, and was the only difference in optimized and baseline temperature setpoint trajectories. The optimized strategies also continued to bring in large amounts of OA until about 9 a.m., slightly later than the economizing of B2, since doing so improved IAQ at only a marginal energy cost. However, the IAQ impact of the longer morning flush was not large, since TVOC concentrations were also low by about 8 a.m. under B2. The other interesting feature of the optimized strategies in the summer, especially those with low values of c, was relatively low ventilation rates during much of the day, interspersed with short periods of higher ventilation. This pattern was especially true for $c = 1$. It appears that, under these conditions, such a punctuated approach of delivering OA may be an energy-effective method of providing dilution. It does not seem that this strategy is entirely an artifact of the optimization process, since smoothing the ventilation rate with a moving average or a median filter reduced energy performance.

July perimeter zone results exhibited lower TVOC concentrations than in the core, especially during the early morning, as a result of the higher infiltration rate. In the winter, this meant that a morning

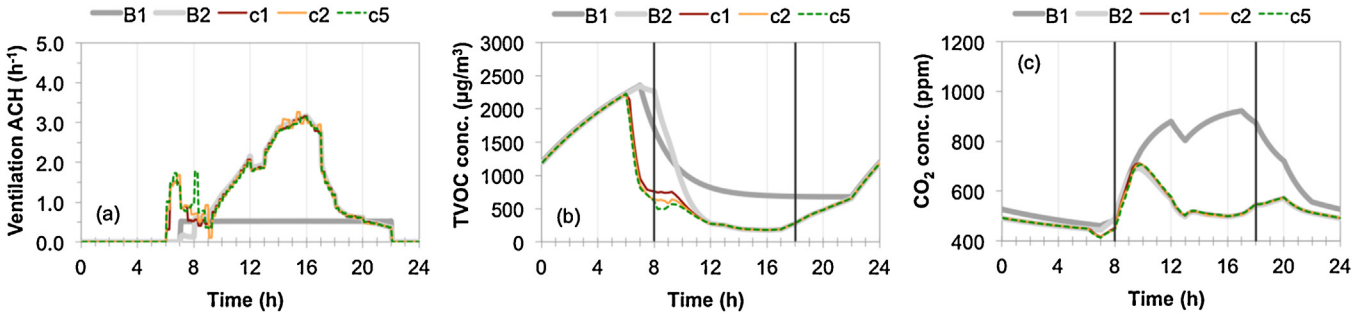

Fig. 4. October results in the core zone: ventilation ACH, TVOC concentration, and CO_2 concentration. For much of the day in all three plots, the curve for B2 is difficult to see because it is directly under those of the optimized strategies.

Table 2
Mean and peak CO_2 and TVOC values for the six cases for two baselines (B1 and B2) and four optimized (opt) strategies, rounded to the nearest multiple of 10. Bolded strategies are discussed in the text.

Case/strategy	Mean CO_2 (ppm)	Peak CO_2 (ppm)	Mean TVOC ($\mu g/m^3$)	Peak TVOC ($\mu g/m^3$)
January, Core zone				
B1	830	930	880	1630
B2	840	890	990	2380
Opt $c=0$	950	1070	1320	2050
Opt $c=1$	940	1060	1160	1600
Opt $c=2$	920	1020	990	1200
Opt $c=5$	850	920	750	800
January, Perimeter zone				
B1	760	830	600	780
B2	750	800	620	1000
Opt $c=0$	850	1020	920	1330
Opt $c=1$	840	1010	890	1300
Opt $c=2$	820	1000	840	1230
Opt $c=5$	790	990	690	860
July, Core zone				
B1	830	920	870	1650
B2	800	890	560	700
Opt $c=0$	1500	2390	1420	2520
Opt $c=1$	820	990	580	800
Opt $c=2$	790	890	550	690
Opt $c=5$	760	830	500	580
July, Perimeter zone				
B1	760	840	610	800
B2	740	810	490	600
Opt $c=0$	1200	1660	1050	1630
Opt $c=1$	820	1000	590	830
Opt $c=2$	780	890	530	680
Opt $c=5$	760	840	500	590
October, Core zone				
B1	820	920	880	1660
B2	560	690	530	2190
Opt $c=0$	570	730	490	1520
Opt $c=1$	560	710	350	750
Opt $c=2$	560	710	330	640
Opt $c=5$	560	710	310	570
October, Perimeter zone				
B1	760	830	610	810
B2	580	720	410	1020
Opt $c=0$	610	820	430	890
Opt $c=1$	590	760	380	660
Opt $c=2$	590	750	370	610
Opt $c=5$	590	750	350	550

flush was justified in the core zone for its IAQ benefits, but not in the perimeter. In the summer, however, very high OA rates in the morning were beneficial from a pure energy perspective in both zones, and both zones had low TVOC concentrations by about 8 a.m. In July, the optimized strategies also did not reduce ventilation significantly earlier in the afternoon in the perimeter than the core, as they did in the winter. In short, although the July nighttime perimeter and core zone TVOC profiles were different, their profiles during the daytime were quite similar in terms of ventilation rates, cooling loads, and especially contaminant concentrations.

The *October core zone* optimized strategies were also similar to B2. Indeed, for most of the day all three shown optimized strategies discovered the economizer OA control exactly (Fig. 4a), demonstrating the near-optimality of economizing under mild weather conditions, when IAQ and energy objectives are not in conflict. However, there was one important difference from B2: the optimized strategies provided significant OA from about 6 a.m. to 9 a.m., before traditional economizing was activated. Economizing is typically only used when a zone would otherwise be warmer than the cooling setpoint (24 °C in this case). The optimized strategies, on the other hand, set the zone temperature setpoints near the *heating* setpoint (21 °C). Given a somewhat chilly morning temperature (10 °C at 6 a.m.) and low internal gains early in the morning, this setpoint trajectory allowed a much higher ventilation rate with almost no energy penalty. The greater ventilation was quite useful in driving down the morning TVOC concentration, as is evident by comparing B2 and the optimized strategies between 8 a.m. and noon (Fig. 4b). Because B2 and the optimized strategies used negligible coil energy in October, the CO_2 profiles are shown in Fig. 4c. They demonstrate that even when the optimized strategies and B2 had different TVOC outcomes, their CO_2 profiles were similar—in this case, identical.

October perimeter results were, once again, similar to the corresponding core results with the differences caused by greater perimeter infiltration. In particular, although the perimeter optimized strategies also reduced the zone temperature setpoints and brought in OA early in the morning, the TVOC concentrations were not as high and so the reductions were not as dramatic.

3.2. Objective values and Pareto frontiers

The IAQ metrics that resulted from these transient profiles are shown in Table 2. These metric values will be considered alongside the energy results by means of Pareto curves (Fig. 5), which plot the IAQ objective of the four IAQ metrics combined according to Eq. (6) versus the energy objective of the extrapolated EUI calculated by Eq. (4). These curves graphically depict the tradeoffs between IAQ and energy. Each c value leads to a point on the curve that represents a different optimal tradeoff between the objectives.

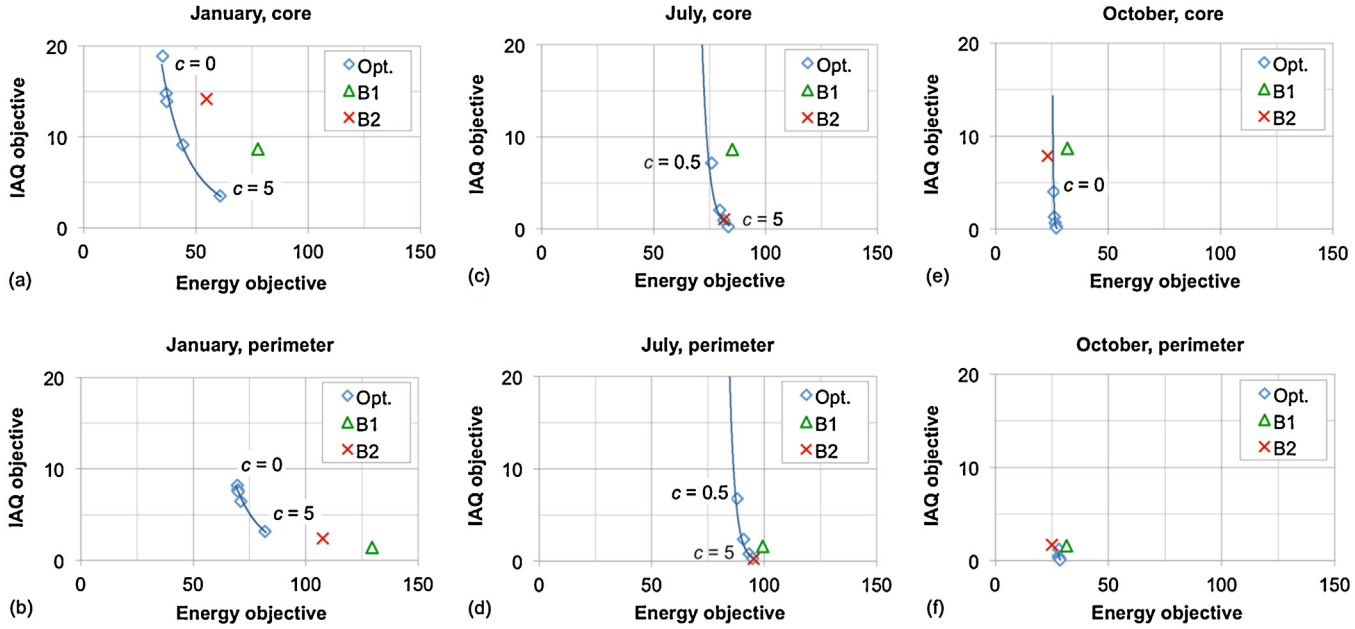

Fig. 5. Pareto curves for the six modeled cases. In July, the points for B2 are difficult to see because they overlap the Pareto curves; they are at the bottom, near the labels for $c = 5$.

Reading from upper left to lower right, the c values in the plots are 0, 0.5, 1, 2, and 5. The two baseline strategies are also plotted. An operational strategy is suboptimal or dominated if it leads to an objective pair (J_1, J_2) to the right of the Pareto frontier. In theory, any pair (J_1, J_2) to the left of the curve is impossible. (In Fig. 5e, B2 is actually slightly to the left, due to the additional hour of fan runtime required of the optimized strategies but not of the baselines.) Again, unlike the energy objective, the IAQ cost is not a real quantity with obvious physical or economic meanings, and should be interpreted by noting the individual metric values in Table 2. In the following discussion, all percent changes are relative to B1 unless otherwise noted.

In *January*, we found the most striking results. In the core zone, B1 was far from optimal (Fig. 5a). Its large IAQ objective value was due primarily to the high peak TVOC that resulted from the nighttime buildup. Using the strategy of economizing-with-DCV (B2) saved 29% of HVAC site energy but significantly degraded IAQ by further increasing the TVOC values. In contrast, the optimized strategy with $c = 2$ saved 43% of HVAC site energy and left the IAQ objective essentially unchanged. Its morning flush and ventilation time-shifting allowed the CO_2 metrics and mean TVOC to rise somewhat, but greatly decreased peak TVOC. A different tradeoff point was provided by the optimized strategy with $c = 5$, which left the CO_2 metric values unchanged, reduced mean TVOC, halved peak TVOC, and still reduced energy use by 22%.

In January in the perimeter zone, both baselines had better IAQ because of infiltration. They also used much more energy than in the core, also in part due to infiltration (Fig. 5b). Unlike in the core zone, using DCV in the perimeter did not degrade IAQ significantly, but it also saved less energy with a 17% reduction. The optimized strategy with $c = 2$, however, reduced HVAC energy consumption by 46%. It did increase all IAQ metric values, but the increases still left the IAQ objective better than that of, for example, the core zone in January under the baselines. Furthermore, the optimized strategy with $c = 5$ changed the IAQ metrics relatively little even as it resulted in 37% energy savings. Thinking about this conversely, one might say that, compared to the optimized strategy with $c = 5$, both of the baseline strategies unintentionally paid very high incremental energy costs for relatively small IAQ improvements.

In *July*, the IAQ metric values under B1 were the same, up to simulation precision, as for the same zones in January. In the July cases, however, B2 used both the economizer and DCV components of its control logic, using the former in the morning and the latter in the afternoon. In both zones, this control only led to 4% HVAC site energy savings. The economizing had the ancillary but more important effect of reducing TVOC levels and, especially in the core, significantly improving the IAQ objective (Fig. 5c). Because introducing large amounts of OA in the morning was favorable from a pure energy standpoint, the optimized strategies had little room to improve. The tradeoffs ranged from saving about 11% of energy and providing worse IAQ than B2 (with $c = 0.5$) to providing essentially identical results to B2 (with $c = 2$ in the core and $c = 5$ in the perimeter).

The fundamental difference between July and January was that in January the constraining factor was the TVOC concentration at the very beginning of the occupied period, which could be reduced with a low-cost flush prior to occupancy. In July, the constraining factor was the TVOC concentration for the entire afternoon, a situation no dynamic strategy effectively addressed. (The only exception might be the punctuated ventilation witnessed with the lower c values, if that effect is confirmed.) In short, the Pareto curve was steep—signifying large IAQ degradation for small energy savings—and the baselines were already close to it.

In *October*, the weather was mild and economizing was favorable nearly the entire day. Moving from B1 to B2 did save 27% of energy in the core and 20% in the perimeter, but the absolute savings were essentially negligible (Fig. 5e and f). Perhaps more importantly, it also reduced every IAQ metric value except peak TVOC, which it increased because of its DCV control. The optimized strategies, on the other hand, found that providing more OA from about 6 a.m. to 9 a.m. had more IAQ benefit than energy cost. For example, compared with B2, the optimized strategy with $c = 1$ saved less energy (17% in the core and 10% in the perimeter) but dramatically lowered the TVOC metric values, reducing peak TVOC by more than 300 $\mu g/m^3$ in the perimeter and 1400 $\mu g/m^3$ in the core. In this case, the Pareto curve was also steep, but here since neither baseline was at the good-IAQ end of it, there were large IAQ improvements available for small energy costs.

3.3. Optimization novelty and approach

These benefits and effects can be attributed to four components of the approach taken here. They are worth distinguishing explicitly, because they influenced the results in different ways. *Optimization based on both energy and IAQ metrics over the day* was the most crucial and novel part of the approach. It was more flexible than either a constant ventilation rate or classical control with a constant concentration setpoint. This flexibility allowed exploitation of dynamic effects, like providing more OA to the core zone early in the morning in January and October. *The availability of perfect information* also accounted for a significant portion of the benefits by enabling more situation-specific control, like introducing less ventilation in the perimeter zone because of its greater infiltration rate. This outcome suggests that there could be a great deal of value to gathering better IAQ information and acting intelligently on it. *The Pareto framework* was useful for exploring tradeoffs, but it cannot indicate what *c* value should be used. *The particular energy and IAQ objective functions* were simple choices for this exploratory study, but further work is needed to develop objectives based on real costs and the best available information about contaminant health, productivity, and comfort impacts. The IAQ objective can easily be modified in the future without altering the overall approach.

4. Conclusions

In January, significant energy savings, from 20% to 45% of HVAC site energy depending on desired tradeoff, were possible. Savings resulted from dynamic strategies such as: a strong flush of ventilation air combined with low zone temperature setpoint early in the morning oftentimes prior to occupancy, but only when a nighttime contaminant buildup made this beneficial; a time-shift of occupied-morning ventilation to mid-day, when OA was less costly or free to condition; a reduction in ventilation late in the afternoon when higher rates would require heating but not be effective at improving IAQ metrics; and a balance of concerns in the perimeter that traded a slight reduction in IAQ for a large savings in energy.

In July and October, energy could be saved compared to fixed ventilation, but not compared to a combination of existing strategies, namely integrated economizing and DCV. One important implication is that these off-the-shelf technologies are near-optimal at many times of the year, if they can be implemented correctly. However, these technologies could be improved upon in October, when significant IAQ benefits at almost no energy cost were achieved by bringing in more OA earlier in the morning. The results so far clearly demonstrate that, under some conditions, the conventional ventilation methods are suboptimal and that the potential benefits of adaptive strategies are large. Knowledge gained from these multiobjective optimizations can help formulate zone-, time-of-day-, and season-specific control rules to complement economizer and DCV operation. For example, prescribing a morning OA flush in a core zone but not a perimeter zone would capture a portion of the benefits observed in January. Simple rules could also be developed to alter ventilation rates based on the amount of infiltration, similar to rules that have been proposed for residences [31].

Determining these and other recommendations will require extending this work by considering other climate zones and weather data, building and system types, objective functions, optimization methods, and models. Exploring the impact of occupant VOC emissions and higher concentrations of outdoor contaminants should also be priorities. As we move forward, thinking about IAQ and energy goals at the same time provides a framework for these investigations.

Acknowledgments

This work was funded by the Energy Efficient Buildings Hub (EEB Hub), located in Philadelphia, PA, funded by the US DOE through the Pennsylvania State University.

References

[1] ASHRAE, Ventilation for Acceptable Indoor Air Quality, American Society of Heating, Refrigerating, and Air-Conditioning Engineers, Atlanta, GA, 2010, ASHRAE 62.1-2010.
[2] Annual Energy Review, US Energy Information Administration September 27 (2012), http://www.eia.gov/totalenergy/data/annual/showtext.cfm?t=ptb0211. Accessed May 16, 2013.
[3] X. Xu, S. Wang, Z. Sun, F. Xiao, A model-based optimal ventilation control strategy of multi-zone VAV air-conditioning systems, Applied Thermal Engineering 29 (2009) 91–104.
[4] J. Laverge, A. Janssens, Optimization of design flow rates and component sizing for residential ventilation, Building and Environment 65 (2013) 81–89.
[5] M.H. Sherman, I.S. Walker, Meeting residential ventilation standards through dynamic control of ventilation systems, Energy and Buildings 43 (2011) 1904–1912.
[6] S. Wang, Z. Ma, Supervisory and optimal control of building HVAC systems: a review, HVAC&R Research 14 (1) (2008) 3–32.
[7] D.S. Naidu, C.G. Rieger, Advanced control strategies for heating, ventilation, air-conditioning, and refrigeration systems—an overview: part I: hard control, HVAC&R Research 17 (1) (2011) 2–21.
[8] D.S. Naidu, C.G. Rieger, Advanced control strategies for HVACR systems—an overview: part II: soft and fusion control, HVAC&R Research 17 (2) (2011) 144–158.
[9] K.F. Fong, V.I. Hanby, T.T. Chow, HVAC system optimization for energy management by evolutionary programming, Energy and Buildings 38 (3) (2006) 220–231.
[10] J.E. Braun, Reducing energy costs and peak electrical demand through optimal control of building thermal storage, ASHRAE Transactions (2) (1990) 876–888.
[11] G.P. Henze, C. Felsmann, G. Knabe, Evaluation of optimal control for active and passive building thermal storage, International Journal of Thermal Sciences 43 (2) (2004) 173–183.
[12] J.M. House, T.F. Smith, System approach to optimal control for HVAC and building systems, ASHRAE Transactions 101 (2) (1995) 647–660.
[13] S.W. Wang, X.Q. Jin, Model-based optimal control of VAV air-conditioning system using genetic algorithm, Building and Environment 35 (6) (2000) 471–487.
[14] N.N. Kota, J.M. House, S. Arora, T.F. Smith, Optimal control of HVAC systems using DDP and NLP techniques, Optimal Control Applications and Methods 17 (1) (1996) 71–79.
[15] G.P. Henze, An overview of optimal control for central cooling plants with ice thermal energy storage, Journal of Solar Energy Engineering—Transactions of the ASME 125 (3) (2003) 302–309.
[16] B. Coffey, F. Haghighat, E. Morofsky, E. Kutrowski, A software framework for model predictive control with GenOpt, Energy and Buildings 42 (7) (2010) 1084–1092.
[17] R. Yang, L. Wang, Optimal control strategy for HVAC system in building energy management, in: Paper Presented at Proceedings of the IEEE Power Engineering Society Transmission and Distribution Conference, 2012.
[18] K. Keeney, J. Braun, Simplified method for determining optimal cooling control strategies for thermal storage in building mass, HVAC&R Research 2 (1) (1996) 59–78.
[19] J.E. Braun, Z. Zhong, Development and evaluation of a night ventilation precooling algorithm, HVAC&R Research 11 (3) (2005) 433–458.
[20] G.P. Henze, A.R. Florita, M.J. Brandemuehl, C. Felsmann, H. Cheng, Advances in near-optimal control of passive building thermal storage, Journal of Solar Energy Engineering—Transaction of the ASME. 132 (2) (2010) 1–9.
[21] U. Satish, M.J. Mendell, K. Shekhar, et al., Is CO_2 an indoor pollutant? Direct effects of low-to-moderate CO_2 concentrations on human decision-making performance, Environmental Health Perspectives 120 (12) (2012) 1671–1677.
[22] T. Salthammer, Critical evaluation of approaches in setting indoor air quality guidelines and reference values, Chemosphere 82 (11) (2011) 1507–1517.
[23] US DOE, EnergyPlus Energy Simulation Software (2012), http://www.eere.energy.gov/buildings/energyplus/
[24] M. Deru, K. Field, D. Studer, et al., US Department of Energy Commercial Reference Building Models of the National Building Stock, National Renewable Energy Laboratory, Golden, CO, 2011, NREL/TP-5500-14686.
[25] B. Griffith, N. Long, P. Torcellini, R. Judkof, D. Crawley, J. Ryan, Methodology for Modeling Building Energy Performance across the Commercial Sector, National Renewable Energy Laboratory, Golden, CO, 2008, NREL/TP-550-41956.
[26] W.R. Chan, Assessing the Effectiveness of Shelter-in-Place as an Emergency Response to Large-Scale Outdoor Chemical Releases [Dissertation], Lawrence Berkeley National Laboratory, University of California, Berkeley, CA, 2006.
[27] A. Rackes, M.S. Waring, Modeling impacts of dynamic ventilation strategies on indoor air quality of offices in six US cities, Building and Environment 60 (2013) 243–253.

[28] M. Wetter, GenOpt Generic Optimization Program, Lawrence Berkeley National Laboratory, Berkeley, 2011.
[29] T.G. Kolda, R.M. Lewis, V. Torczon, Optimization by direct search: new perspectives on some classical and modern methods, SIAM Review 45 (3) (2003) 385–482.
[30] E. Polak, M. Wetter, Precision control for generalized pattern search algorithms with adaptive precision function evaluations, SIAM Journal on Optimization 16 (3) (2006) 650–669.
[31] W.J.N. Turner, M.H. Sherman, I.S. Walker, Infiltration as ventilation: weather-induced dilution, HVAC&R Research 18 (2012) 1122–1135.