SimBuild 2006

Second National IBPSA-USA Conference
Cambridge, MA
August 2-4, 2006

A MULTIPLE-BUILDING OPTIMIZATION SCHEME BASED ON STATISTICAL BUILDING-LOAD MODELS

Leslie K. Norford[1] and Hai-Yun Xing[1,2]
Building Technology Program[1]
Department of Architecture
Massachusetts Institute of Technology
Cambridge, MA 02139
Credit Suisse, New York, NY[2]

ABSTRACT

This paper presents an optimization process for multiple-building load aggregation based on statistical models of building dynamics. The aggregation process is modeled as a central optimizer and participating buildings as unit simulators. By learning from whole-building simulation results, unit simulators describe individual buildings' base-load profiles and responses to control parameter changes. The optimizer makes supervisory control decisions based on a centralized objective function which takes inputs from individual simulators. A non-linear mathematical optimization problem is formulated and solved using an interior-point-method commercial solver. Results are compared to simple cases for which optimal or near-optimal results are relatively easy to find. The expansion of the analytical framework is discussed.

INTRODUCTION

An effective means of maximizing savings in today's power markets, load aggregation is the process by which individual energy users band together in an alliance to secure more competitive prices than they might otherwise receive working independently. Aggregation can be accomplished through a simple pooling arrangement or through the formation of clusters where individual contracts are negotiated between the suppliers and each member of the aggregate group. Aggregation offers load diversity among multiple facilities, improves load factors, and leads to a smaller demand charge.

The diversified peak load can be further reduced via active control measures. These measures include those that shift peak-period loads to off-peak periods and those that curtail loads on a quick-response basis. Load control for an aggregation of buildings is more complex than simply minimizing peak loads in individual buildings because the aggregated peak may be dominated by a single building, making it necessary to reduce loads in other buildings not at the time of their peak but at the time of the aggregated peak. However, aggregated load control offers an opportunity for load-control schemes that are simpler to implement than those often used for a single building. For example, instead of carefully controlling thermostat set points over time in all buildings, it may be possible to simply turn off air conditioners or chillers for short periods, timing these events in a way that minimizes the aggregated peak.

Regardless of the form of control, models of building response are important. Most previous studies have focused on a single building. Braun et al. (2001) developed and used an inverse model to explore the effect on energy cost of different building-thermal-mass control strategies. The inverse model was based on a transfer function and used measured data to 'learn' system behavior and provide relatively accurate site-specific performance predictions. Component (fan and chiller) power models were quadratic functions of flow or temperature variables. Wright and Farmani (2001) simultaneously optimized the building fabric and HVAC supervisory control strategies using a genetic algorithm. A lumped-capacitance model was used to represent the thermal response of the zone, while the HVAC system performance was simulated using steady component models.

Norford and Reddy (2004) examined short-term load-control measures in a single building and in pairs of buildings. They used a nodal thermal model that lacked the detail of more comprehensive models, a genetic algorithm to optimize load control for one building, and limited combinations of load control measures for two buildings. Xing (2004) extended their work by using detailed simulations (EnergyPlus 2006) of larger sets of load-control options for larger groups of buildings. To reduce the combinations of options for multiple buildings (100 possible hourly thermostat set point adjustments over a peak-load period for one building yields 10^{10} possibilities for a five-building aggregation), she developed a set of rules to narrow the search to feasible options that did not violate comfort constraints. Her enumeration method was not exhaustive but what she termed "smart;" for example, if a small thermostat increase led to a violation of thermal-comfort criteria, larger increases in set point were not evaluated. The enumeration was also restricted; in particular, thermostat increases were held constant during the load-reduction period, which varied from 1-4 hours.

Xing also applied a genetic algorithm (GA) to single-building and multi-building load control, using a publicly available GA toolbox (GAOT 2006). Notably, as shown in Figure 1, the single-building GA produced hourly thermostat set points that approximated the exponential trajectories determined as optimal by Braun and Lee (2006).

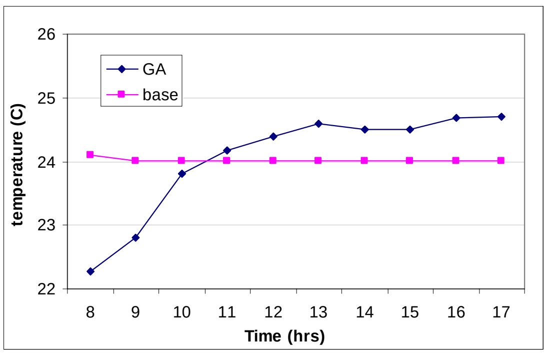

Figure 1. Optimal hourly thermostat set points to minimize daily electrical energy and demand costs.

For multiple buildings, Xing applied her enumerative search and the GA to short-term load control and to fan-driven night cooling. For five-building aggregates, both approaches involved a large number of simulations and the optimization was slow. A detailed first-principles or data-driven model is useful in single-building load optimization, but is perhaps not the best choice for a multiple-building problem. In order for the central optimizer to make reasonably quick decisions for all participating buildings, the number of parameters is crucial.

This paper, which summarizes work in Xing (2004), presents a new way of optimizing the multi-building load aggregation process, based on statistical building-load models that are learned from whole-building simulations. It first presents a mathematical formulation for the problem and then establishes a time-series model to predict load profiles for base-case buildings. Load reduction via short-term increases in temperature set points and the associated degradation in thermal comfort for the base-case buildings are modeled through linear regression. The multi-building aggregation is then formulated as a nonlinear optimization problem supported by the individual load-prediction models. Optimization results are evaluated and compared to those from enumerated simulations and from the GA.

PROBLEM FORMULATION

The general problem is to minimize the aggregated total electricity cost or the peak demand while maintaining a certain comfort level in all of the aggregated buildings. We have represented hourly power consumption and the thermal-comfort index, predicted mean vote (PMV), in each individual building as functions of building dynamics, load control parameters and time. Equation 1 governs the minimization of peak load for an N-building pool, with comfort as a constraint. The method developed for the peak-load problem can be applied to the total-cost problem without any extra complexity, because the sum of hourly energy costs is a simple linear term.

$$\begin{aligned}
&\min \quad \max_t \left(\sum_{i=1}^{N} W_{i,t} \right) && \text{Minimize the peak demand} \\
&s.t. \quad W_{i,t} = f_i\left(x_{i,t}, x_{i,t-1}, \ldots, u_{i,t}\right), \forall i && \text{electricity consumption as a function of building dynamics} \\
&\qquad PMV_{i,t} = g_i\left(x_{i,t}, u_{i,t}\right) \le PMV_i, \forall i && \text{comfort requirement} \\
&\qquad u_- \le u_{c,i,t} \le u_+, \ c \in C, \forall t, \forall i && \text{continuous control variables} \\
&\qquad u_{d,j,t} = 0 \ or \ 1, \ d \in D, \forall t, \forall i && \text{discrete control variables}
\end{aligned} \tag{1}$$

where

$W_{i,t}$	=	electricity consumption by building i at time t, i = 1,…N, t = 1,…24
$PMV_{i,t}$	=	PMV values in building i at time t
PMV_i	=	PMV requirement in building i
D	=	demand charge rate
$u_{i,t}$	=	continuous, $u_{r,i,,t}$, and discrete, $u_{z,i,,t}$, control variables
$x_{i,t}$	=	state variables

$f_{i,}$ = building i dynamics determining electricity consumption at hour t
g_i = building i dynamics determining thermal comfort i is at hour t

This is a min-max problem: the program first looks for the peak demand, which is a maximization problem over a 24-hour period on the aggregated load profile, then minimizes this peak demand over the control variables specified in both equations. Min-max problems are difficult to solve in general but can be converted to a typical minimization problem by replacing the maximization step with 24 inequalities, one for each hour, in the constraint. In this case, the hourly loads for each building are adjusted to minimize their sum during the peak-load hour.

BASE LOAD PREDICTOR – A TIME SERIES MODEL

The base building was a three-zone space on the top floor of a commercial building with a variable-air-volume ventilation system. Load profiles, generated in EnergyPlus (2006) were varied by adjusting thermal mass and window size and orientation to create multiple "buildings" for the aggregation. Figure 2 shows the load profiles.

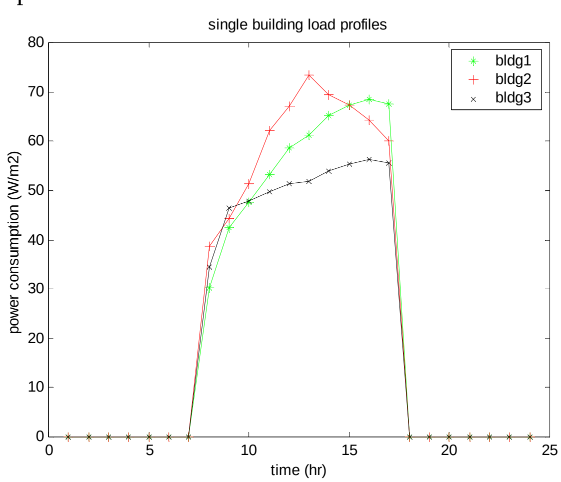

Figure 2. Load profiles for the three test buildings.

We limited this study to an exploration of the benefits of short-term changes in temperature set points, as would be required after notification of an upcoming utility peak-load period. On-off switching of equipment was excluded. Night pre-cooling was considered only for two-building aggregates, using enumeration and a GA. The limitation of a four-hour peak-load period without pre-cooling restricted the number of simulations needed to demonstrate the concept. For each of the individual buildings, we ran simulations from June 1 to August 31 using Los Angeles weather data for nine scenarios: a base case without temperature set-point change, set-point increases of 1-4 ^{o}C from hours 13-17, and a 2 ^{o}C set-point increase over 1-4 hour periods. June data were used for training and July and August for testing. The base-case data were sufficient for training the base-load model; the remainder was used to train load-reduction models. Internal loads were limited to a single daily schedule. Hourly outputs from the simulations were total power, PMV values, and indoor air temperatures.

We chose a time-series model to capture the building dynamics because such models use a small number of parameters, key to a large multi-building problem, and their linearity makes it easy to solve a nonlinear optimization problem for which the time series model is a constraint. All the data available from the simulations were normalized before regression, which was done with a generalized linear model using Iterative Reweighted Least Squares (IRLS) (S-Plus 2006). Autocorrelations applied to a time series of total power showed that power at times lagged by one, 24 and 25 hours were important predictors of current power. Total power and exogenous variables at the current hour and the three lagged times were therefore included in the regression; exogenous variables were outdoor temperature and solar loads on the roof and the east, south and west facades. From the regression, the power terms dominated the prediction, as determined by the t statistic; the influence of the exogenous variables was implicit in the power but explicitly small and we limited the model to the power terms. The lack of one-hour-lagged power measurements for a prediction of power beyond a single hour, as is needed to control load over an upcoming peak period, was compensated by using predicted peak daily values of outdoor temperature (following Seem and Braun 1991) and solar radiation to predict daily peak power, and scaling the hourly power prediction with the estimate of the peak power. Peak-power correction was especially important on two of the four days shown in Figure 3.

LOAD-REDUCTION MODEL

Simulated increases in thermostat set points led to load reductions that varied linearly with base load for a given hour; Figure 4 shows a representative relationship. The set point was held constant during the load reduction period; while this is not optimal for a single building (see Figure 1), constant set-point changes in individual buildings can be coordinated across a group to affect the aggregate peak. Load reductions varied approximately with the square root of the increase in thermostat set point, as shown in Figure 5; the relationship is not linear because temperature does not respond instantly to changes in set point. From these observations, we developed the load-reduction model shown in Equation 3.

$$\Delta P = b_{i,0} + b_{i,1}(\Delta T)^{0.5} + b_{i,2}B \qquad (3)$$

where

ΔP	=	power reduction from the base case
ΔT	=	thermostat set point increase
B	=	base load

After establishing values for the model coefficients by regression, we compared predicted decreases in load with those calculated directly from simulation. Figure 6 shows the results; the agreement is relatively poor for a 1 ^{o}C thermostat increase but approaches the desired 45^{o} line for larger increases.

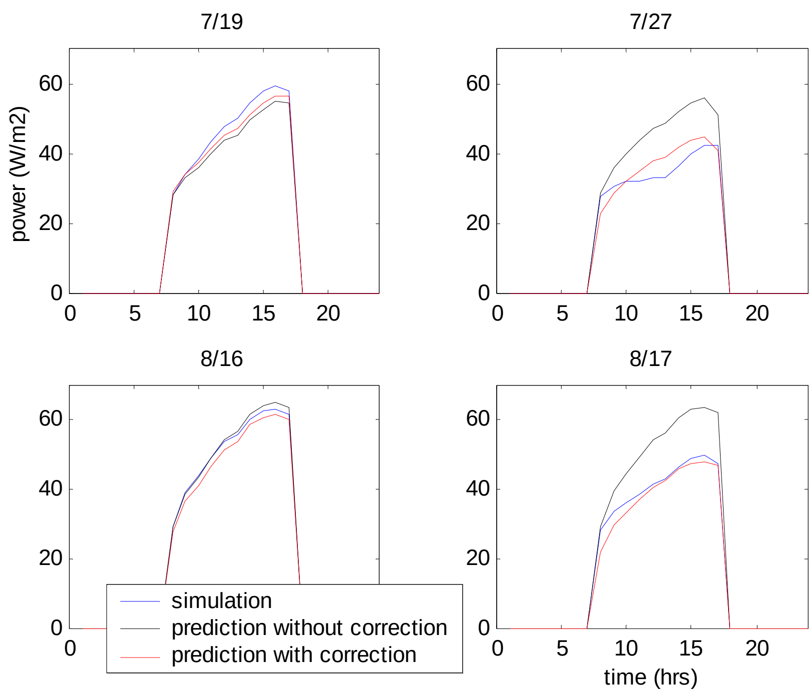

Figure 3. Power prediction, with and without peak-power correction.

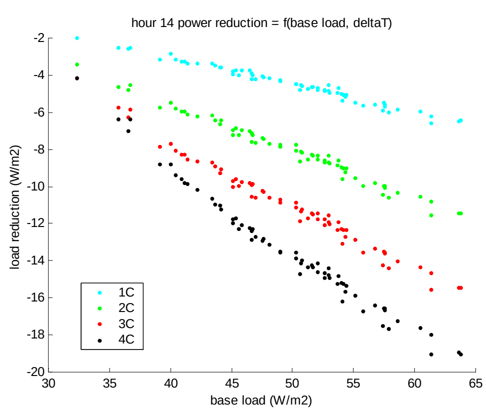

Figure 4. Load reduction from the base case at hour 14, as a function of base load and thermostat change.

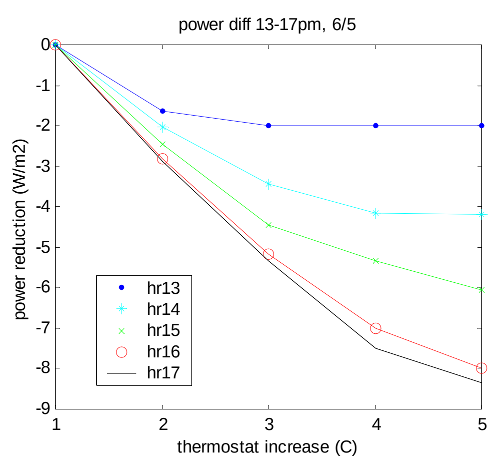

Figure 5. Power difference from the base case at hours 13-17 as a function of thermostat change.

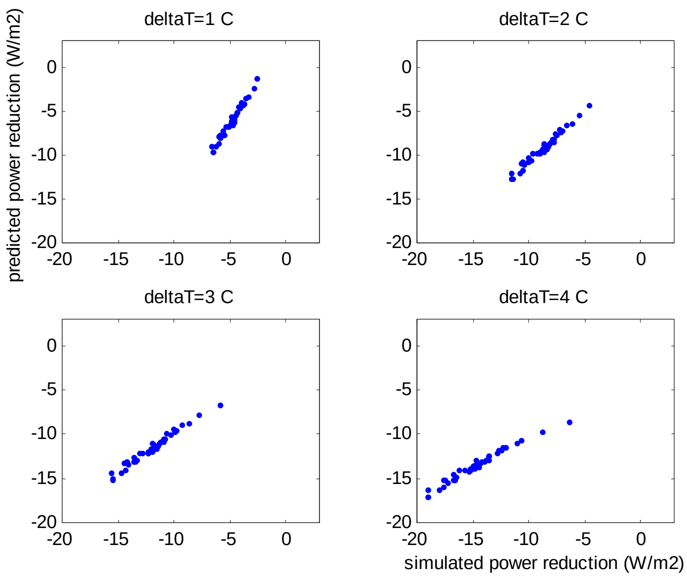

Figure 6. Predicted power reduction prediction vs. simulation for hour 14.

Modeling power reductions for a general load-control strategy is more complicated than for thermostat control. One key issue is interdependence among hours. For example, if the thermostat set point is increased at hour 14 alone, hour 15 will certainly see a power spike associated with returning the temperature to a lower value and hours 16 and 17 will experience a small amount of power increase as well. The challenge is how to represent the large number of possible schedule combinations, especially when a longer control period is involved. Integer programming might be used to assign an integer variable to each hour indicating whether this particular hour incurs a set-point-recovery spike. Then we could run a regression to find the response for a specific hour as a function of a variety of inputs. We ignore the temperature-recovery issue in this model by applying the load shedding, regardless of the magnitude, to all the afternoon hours. We understand that we are trading optimality for simplicity: the constraint we apply to the process could cost us better solutions.

COMFORT MODEL

The comfort model estimates how much the thermostat set points can rise during the peak-load period without leading to indoor temperatures that exceed a constraint on thermal comfort. The upper bound for thermal discomfort was a PMV of 1.5, which corresponds to 50% dissatisfied occupants. The predicted maximum PMV for the next day was derived from a linear relationship between peak load and PMV, shown in Figure 7, and a prediction of next-day peak load in the absence of control measures. For example, the daytime-peak PMV is predicted to be 1.2 when the peak load is predicted to be 55 W/m^2. In this case, PMV could increase by 0.3 before reaching the limit. PMV increases relative to the base case were observed to vary nearly linearly with thermostat increase for a given hour, as shown in Figure 8. However, the influence of hour was sufficiently small to be excluded from the model. The PMV increase was also found to be independent of base load over a wide range of base loads. The allowable increase in PMV was therefore mapped to a uniform increase in thermostat set point for the load-control period via a linear regression.

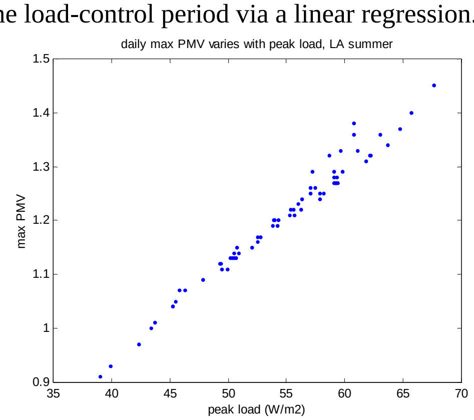

Figure 7. Relationship between daytime peak PMV values and peak load in summer.

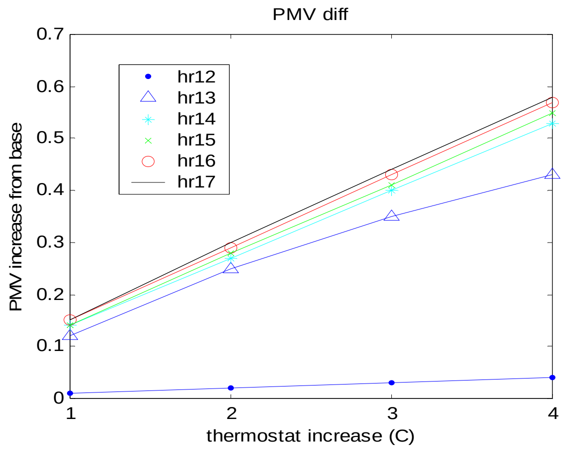

Figure 8. PMV increase at hours 12-17 as a function of thermostat change, for a single day.

NONLINEAR CENTRAL OPTIMIZER

The final model, Equation 4, incorporates the base-load predictor, peak-load corrector, load-reduction model and the PMV-increase model. This model is fairly easy for a commercial nonlinear solver to handle. The cost function and most constraints are linear, and the only nonlinear constraint is quadratic, a "soft" and computationally friendly one. We coded this model in a platform for optimization problems with complicated constraints (AMPL 2006) that is supported by a variety of linear and nonlinear solvers. This was an attractive alternative to coding the problem from scratch. The optimizer we selected (Vanderbei 1997) uses an infeasible primal-dual interior-point method applied to a sequence of quadratic approximations to the given problem. The problems can be linear or non-linear, convex or non-convex, constrained or unconstrained. The only real restriction is that the functions defining the problem be smooth, i.e., they should be twice continuously differentiable at every point visited by the algorithm.

$$
\begin{aligned}
&\min \quad z && \text{total peak load} \\
&s.t. \quad \sum_{i=1}^{N} W'_{i,t} \le z, \forall t && \text{together with the cost function to handle a min max problem,} \\
& && W'_{i,t} \text{ and } W_{i,t} - \text{the real and base power use of building } i \text{ at time } t \\
& && \qquad (4) \\
&W'_{i,t} = \left(a_{i,1}W_{i,t-1} + a_{i,2}W_{i,t-24} + a_{i,3}W_{i,t-25}\right)\times \alpha_i && \text{linear base-load prediction and peak-load correction} \\
&\qquad + f\left(\left(\Delta T_{i,t}\right)^{0.5}, W_{i,t}\right),\ \forall i, \forall t \ge 2 && \text{nonlinear power reduction} \\
&dPMV_{i,t} = g\left(\Delta T_{i,t}\right) \le \Delta PMV_i,\ \ \forall i, t \in [14,17] && \text{PMV prediction and comfort requirement} \\
&\Delta T_{i,t} \in \left[\Delta T_{low}, \Delta T_{high}\right], \forall t, \forall i && \text{thermostat set point range}
\end{aligned}
$$

RESULTS

Table 1 compares the performance of the enumerative search and the GA for pairs of buildings. The GA produced larger peak reductions due to the non-exhaustive and therefore sub-optimal nature of the enumerated simulations but required more time, especially for night cooling, when thermostat adjustments were considered for all occupied hours. The enumerative search took less than 0.1 minute to match the single-building simulations and select the optimum, using a general-purpose programming package (Mathworks 2006). However, this process scales as $O(N^n)$, where N is the number of feasible solutions and n is the number of buildings. For five buildings, a mix of types 1 and 2, the enumeration required 150 minutes for the matching in addition to 12 minutes for the simulations. The GA required 840 simulations and 140 minutes for the same five-building aggregation and 1236 simulations for a five-building aggregation consisting of types 1 and 3; when compared with the two-building cases in Table 1, both five-building examples indicated that the required time for the GA scaled linearly with the number of buildings.

Table 1. Optimizing two-building thermostat set points with enumeration and a GA.

Load control and buildings		peak W/m^2	peak time hour	peak reduction	simulations	Time min
5-hour peak; bldg 1 and bldg 2	Base	135	13			
	Enum	128	13	5.2%	64	12
	GA	126	15	6.7%	340	57
5-hour peak; bldg 1 and bldg 3	Base	125	16			
	Enum	112	13	10.4%	107	20
	GA	111	14	11.2%	536	90
5-hour peak; bldg 2 and bldg 3	Base	125	13			
	Enum	119	12	4.8%	107	20
	GA	118	14	5.6%	514	86
Night cool; bldg 1 and bldg 2	Base	137	13			
	Enum	123	15	8.9%	64	12
	GA	121	15	10.4%	1163	195

To evaluate the statistical-model-based, multi-building optimizer (SMBO), we first ran the program for a single building without comfort constraints as a test case: total demand as the cost function and hour 14-17 thermostat set points as control variables. With a starting point of no thermostat increases, the code took essentially no time to get to the global optimum of 4 ^{o}C temperature set-point increases in hours 14-17. We then ran the program for two two-building cases: two identical buildings of model type 1; and one building each of types 1 and 3. Figure 9 compares four profiles for the case with two type-1 buildings: the sum of two base loads from simulation, without any optimization; the optimal two-building aggregate from enumeration; the sum of two base loads as predicted by the statistical models; and the aggregated load-control results from the SMBO. The SMBO's predictions of base and aggregated load profiles differ from the simulation and enumeration results, but the reduction due to aggregated load control is essentially captured, as summarized in Table 2.

For five buildings in the aggregation pool, four of type 1 and one of type 3, the SMBO still took no time to do the optimization after the simulations were completed, compared to two hours for the enumeration method. The savings again were reasonably close but not in excellent agreement: non-exhaustive enumeration produced a 9.2% peak reduction while the SMBO predicted 6.8%, shown in Table 2. Based on Figure 8, the discrepancies are thought to be due primarily to the load models used by the SMBO. Overall, the operations decisions between these the enumeration and SMBO methods were similar, but the latter failed to show the fact that the same models in an aggregation pool can act quite differently. All buildings of type 1 in the SMBO method acted similarly and their thermostat increases were close to each other, but buildings in the enumeration acted on different schedules and therefore achieved larger savings. The SMBO needs to be improved to evaluate more flexible scheduling.

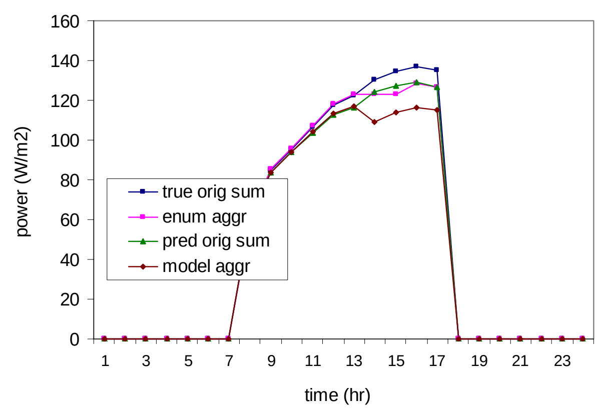

Figure 9. Aggregating two identical buildings of model type 1: statistical-model-based optimizer vs. simulation and enumeration.

Table 2. Statistical-model-based optimizer vs. enumeration

		Peak load (W/ m^2) and hour, before and after load shift	Peak reduction	Hourly thermostat set points and maximum PMV	Computation Intensity
bldg 1 bldg 3	Enumeration	125 / hr 16 -> 112 / hr 12	10.4%	bldg 1 (0,1,1,1) bldg 3 (2,2,2,2) PMV <= 1.5	123 simulations for every optimization
	SMBO	114 / hr 16 -> 104 / hr16	8.6%	bldg 1 (0.5,0.5,0.5,0.5) bldg 3 (3, 3, 3, 3) PMV <= 1.5	460 simulations to prepare training data, one time only; optimization takes no time
bldg 1 bldg 1	Enumeration	137 / hr 16 -> 128 / hr 16	6.6%	1st bldg 1 (1,1,1,1), 2nd bldg 1 (0,1,1,1) PMV <= 1.5	64 simulations for every optimization
	SMBO	128 / hr 16 -> 121 / hr 16	5.7%	both buildings (0.5, 0.5, 0.6, 0.6) PMV <= 1.5	460 simulations to prepare training data
bldg 1 bldg 1 bldg 1 bldg 1 bldg 3	Enumeration	330 / hr16 -> 299 / hr 16	9.2%	(0,1,1,1) (1,1,1,1) (0,1,1,1) (0,0,1,1) (2,2,2,2)	123 simulations plus two hours for optimization
	SMBO	305/ hr16 -> 280 / hr16	6.8%	All 4 bldg 1 (0.6, 0.5, 0.5, 0.5) bldg 3 (2, 2.5, 3, 3) PMV <= 1.5	460 simulations to prepare training data

A 10-building aggregation would be computationally infeasible for the enumeration and formidable for the GA. The SMBO required very little computation time for several 10-building aggregations. To run the model-based optimizer, we need to have either simulation or experimental data to train a unit simulator for each building participating in an aggregation. We only trained models for buildings 1 and 3 in this research.

Table 3 shows the resulting savings. The mix within a pool matters. Equal numbers of building 1 and building 3 achieved 2% more peak reduction than having nine buildings of type 1 and one building of type 3; Individual buildings played different roles as a consequence of their load profiles and thermal mass. The savings for the aggregate tended to favor more buildings of type 3 than type 1.

Table 3. Aggregation using model-based approach

cases	1 bldg1 9 bldg3	3 bldg1 7 bldg3	5 bldg1 5 bldg3	7 bldg1 3 bldg3	9 bldg1 1 bldg3
peak reduction	8.4%	8.6%	8.7%	8.0%	6.5%

CONCLUSION

This paper developed an optimizer for peak loads and total costs of aggregates of buildings using a statistical-model-based approach. A simple time-series model was used to represent building dynamics. Regression was applied to correct the peak load and scale the diurnal power profile. Two regression models were developed to describe load reduction and comfort degradation respectively from the base case due to load control based on increases in thermostat set points. A nonlinear optimization scheme brought these parts together and solved the soft-nonlinear-constraint problem in seconds.

For the peak-load optimization, the statistical approach produced reductions relative to a base case similar to those in an enumerative search or a genetic algorithm, but the base-load and savings predictions were off by a certain amount as a price of simplicity; savings predictions were 14-27% lower than those from enumeration. For a large number of buildings, the statistical-model-based optimizer was computationally efficient and would remain efficient as long as the constraints remain linear or soft nonlinear.

More work is needed in several areas. Discrete control variables should be included in the optimizer, to allow equipment to be turned off. This would require changing the optimization algorithm from an interior-point method to another technique, such as the branch-and-bound method. The load-reduction algorithm should account for load-control actions that produce increases in power during a recovery period. This step is needed to model staged air-conditioner curtailment in a set of buildings. Finally, the stochastic nature of weather and equipment use as influenced by occupant behavior should be assessed.

ACKNOWLEDGMENT

The authors acknowledge the financial support of the California Energy Commission and advice, encouragement and technical assistance from Philip Haves and Michael Wetter.

REFERENCES

AMPL. 2006. A modeling language for mathematical programming. http://ampl.com

Braun, J.E., K. Montgomery and N. Chaturvedi. 2001. "Evaluating the performance of building thermal mass control strategies." *Int. J. of HVAC&R Research* 7(4):403-428.

Braun, J.E. and K.-H. Lee. 2006. "Assessment of demand limiting using building thermal mass in small commercial buildings." *ASHRAE Transactions* 112(1).

EnergyPlus. 2006. EnergyPlus version 1.3. www.eere.energy.gov/buildings/energyplus/ Version 1.1.0 was used in the reported study.

GAOT. 2006. Genetic Algorithm Optimization Toolbox. http://www.ie.ncsu.edu/mirage/GAToolBox/gaot/

Mathworks. 2006. MATLAB 7.2. Natick, MA. www.mathworks.com

Norford, L.K. and T.A. Reddy. 2004. "Tools and capabilities for pro-active multi-building load management: part 2 – aggregated operation." *ASHRAE Transactions* 110(2).

Seem, J.E. and J.E. Braun. 1991. "Adaptive methods for real-time forecasting of building electrical demand." ASHRAE Trans. 97(1).

S-Plus. 2006. Statistical analysis software. Insightful Corp. http://www.insightful.com/.

Vanderbei, R. 1997. "LOQO user's manual, v 3.10." Princeton University.

Wright, J. and R. Farmani. 2001. "The simultaneous optimization of building fabric construction, HVAC system size, and the plant control strategy." *Proc. of Building Simulation 2001*, Rio de Janeiro, Brazil.

Xing, H.-Y. 2004. "Building load control and optimization." Ph.D. thesis, Department of Architecture, Massachusetts Institute of Technology, Cambridge, MA.