General Methodology Combining Engineering Optimization of Primary HVAC&R Plants with Decision Analysis Methods—Part II: Uncertainty and Decision Analysis

Wei Jiang, PhD
Associate Member ASHRAE

T. Agami Reddy, PhD, PE
Member ASHRAE

Patrick Gurian, PhD

Received January 13, 2006; accepted May 9, 2006

A companion paper (Jiang and Reddy 2007) presents a general and computationally efficient methodology for dynamic scheduling and optimal control of complex primary HVAC&R plants using a deterministic engineering optimization approach. The objective of this paper is to complement the previous work by proposing a methodology by which the robustness of the optimal deterministic strategy to various sources of uncertainties can be evaluated against non-optimal but risk averse alternatives within a formal decision analysis framework. This specifically involves performing a sensitivity analysis on the effect of various stochastic factors that impact primary HVAC&R plant optimization, such as the uncertainty in load prediction and the uncertainties associated with various component models of the equipment. This is achieved through Monte Carlo simulations on the deterministic outcome, which allow additional attributes, such as its variability and the probability of insufficient cooling, to be determined along with the minimum operating cost. The entire analysis is then repeated for a specific non-optimal but risk-averse operating strategy. Finally, a formal decision analysis model using linear multi-attribute utility functions is suggested for comparing both these strategies in a framework that explicitly models the risk perception of the plant operator in terms of the three attributes. The methodology is demonstrated using the same illustrative case study as the companion paper.

INTRODUCTION

Incomplete knowledge, as well as noisy or erroneous data, can all be viewed as sources of uncertainty. The existence of uncertainty transforms conventional deterministic process models into stochastic ones, the solution of which indisputably remains challenging yet is of great practical importance. Beginning with the works of Dantzig (1955), Tintner (1955), and Charnes and Cooper (1959), optimization under uncertainty experienced rapid development in both theory and algorithms. There are several approaches for dealing with optimization under uncertainty—sensitivity analysis, stochastic programming, and robust optimization (Mulvey et al. 1995), to name a few. While sensitivity analysis is a reactive approach to controlling uncertainty in that it is used to assess the impact of uncertainty on the optimization result without providing a controlling mechanism, the other two are constructive in that they can yield solutions that are less sensitive to the model data than classical mathematical programming.

The issue of uncertainty, as it relates to the optimization of HVAC&R plants, has been addressed in only a few papers. The effects of many factors affecting the cooling load in a build-

Wei Jiang is a research engineer at Pacific Northwest National Laboratory, Richland, WA. **T. Agami Reddy** is a professor and **Patrick Gurian** is an assistant professor in the Department of Civil, Architectural and Environmental Engineering, Drexel University, Philadelphia, PA.

ing, as well as several parameters of the dynamic chiller sequencing (DCS) procedure described in the companion paper (Jiang and Reddy 2007), have been systematically explored by Olson (1987), who found that the DCS is fairly insensitive to errors in cooling load prediction. Henze and Krarti (1999) investigated the effect of forecasting uncertainty on the cost savings performance of a predictive optimal controller for thermal energy storage. They concluded that the predictive optimal controller is fairly robust and does not require high levels of accuracy in predicting the cooling loads and the real-time pricing (RTP) electricity rates. Henze et al. (2003) investigated whether thermal storage systems can be controlled effectively in situations where cooling loads, non-cooling loads, non-cooling electrical loads, weather information, and the cost of electricity are uncertain and have to be predicted. The analysis concluded that predictive optimal control strategies were greatly superior in reducing utility cost compared to conventional partial-storage thermal energy storage strategies, even with inaccurate forecasts.

To date, most publications addressing optimization and control of HVAC&R plants were based on sensitivity analysis. Further, most of the HVAC&R literature limits itself to load prediction and electrical rate uncertainties (which can be referred to as *external prediction uncertainties*), while the other two categories of uncertainties, model-inherent uncertainty and process-inherent uncertainty, are largely ignored. It could well be that most optimization problems of HVAC&R plants do not warrant formulating them as stochastic problems due to the low degree of associated uncertainty (for example, current load forecast algorithms are relatively accurate, and so are the modeling techniques and control devices). However, studies specifically addressing the aspect of stochastic optimization would be helpful in supporting (or disproving) such a contention.

Furthermore, it may not always be clear beforehand whether the optimal strategy is robust or not and what potential advantage the best stochastic strategy could be expected to provide over the best deterministic strategy (Seferlis and Hrymak 1996). So, it is important to first determine the best deterministic operating strategy and then investigate the additional benefit of using stochastic optimization, which for many systems may be small and may not justify the additional modeling, optimization, and hardware costs.

OBJECTIVES

Despite recent advances in computer power and the development of better optimization algorithms, only few are used in industry. What is more remarkable is that most complex HVAC&R plants are still scheduled by humans in a heuristic manner without the aid of computer supporting tools. One possible reason for this often voiced by professionals is the lack of consideration of how to combine pure engineering solutions with individual risk attitudes of how system operators weigh risk over predicted outcome. The companion paper (Jiang and Reddy 2007) presents a computationally efficient deterministic engineering optimization for scheduling and controlling complex primary HVAC&R plants over a specified planning horizon (taken as 12 hours). It involves developing response surface models for different combinations of system configurations and then using them for static optimization at any given hour. This capability is then used in conjunction with the modified Dijkstra's algorithm (Olson 1987) for dynamic scheduling and optimal control under different operating conditions and pricing signals.

Granted, the operating cost of a plant is a critical criterion or attribute for choosing an appropriate operating strategy; however, it is not the only factor. Different operating strategies may result in different uncertainty profiles with different robustness levels. Minimizing the operating cost and minimizing the risk associated with this optimal strategy are two separate issues altogether; how to trade off between these objectives is basically a problem of decision analysis. Furthermore, different people have different risk attitudes and thus are willing to accept different levels of risk. How to combine different risk preferences into engineering results and, more basically, how to define risk in such a setting are issues investigated in this paper. The objectives of this paper are to

- systematically assess the robustness of the identified optimal deterministic operating strategy when subject to various sources of uncertainties; in other words, to ascertain the extent to which the optimal solution is stable and evaluate the relative effect of different sources of uncertainty on this optimal solution (this is called *sensitivity analysis*);
- compare the uncertainty of the optimal strategy found from our rigorous optimization with those of more conservative non-optimal heuristic strategies;
- model the trade-off between the expected value of the savings, its variability, and the increased probability of being unable to provide sufficient cooling to meet the building load (called *loss of load capability*) and combine these results into a decision analysis framework involving the individual's attitude toward risk and returns; and
- apply this model to both the optimal and heuristic strategies in order to identify the one with the higher operator utility value.

Neglecting the last three aspects is probably one of the reasons why many of the purely analytical optimization solutions are not widely embraced by plant operators. Finally, this paper illustrates optimization methodology and decision analysis techniques using the same hybrid chiller plant selected in the companion paper for two electric price signals, real-time pricing (RTP), and time-of-use rates (TOU) (Jiang and Reddy 2007).

DECISION ANALYSES

Decision analysis provides a formal conceptual framework involving an "overall paradigm and a set of tools by which a decision maker can construct and analyze a model of a decision situation" (Clemen and Reilly 2001). Such methods have been applied in various fields to choose among alternatives in an optimal fashion, including complex engineering systems, economics, physics, social sciences, medical decision making, and others (Hillier and Lieberman 2001). First of all, a decision analysis approach is warranted only when one is faced with a situation involving conflicting objectives, uncertainty in the payoffs, and more than one appropriate alternative. Only once this is determined does one decide on statistical methods for simulating the uncertainty and modeling the results so as to reflect the individual's risk perception or attitudes (low or high risk).

According to Pistikopoulos (1995), a suitable classification of different types of uncertainties is based on the nature of the source of uncertainty: (1) *model-inherent uncertainty*, which includes uncertainties of the various component models arising from inaccurate or incomplete data and/or lack of perfect regression fit in the response model; (2) *process-inherent uncertainty* due to the deadbands within which the decision variables can be controlled in practice (such as the chilled-water temperature control of each chiller and the fan speed control of each cooling tower); and (3) *external prediction uncertainty*, which includes the errors in predicting the system driving functions, such as the building thermal load profiles, wet-bulb temperature profile, and electricity rate profile under RTP. How the various sources of uncertainty appear in the overall framework involving engineering optimization and decision analysis is shown in Figure 1. Several of the modules involving deterministic engineering optimization over a planning horizon are described and illustrated in the companion paper (Jiang and Reddy 2007). The sources of different uncertainties and statistical models accounting for these uncertainties are discussed in the following.

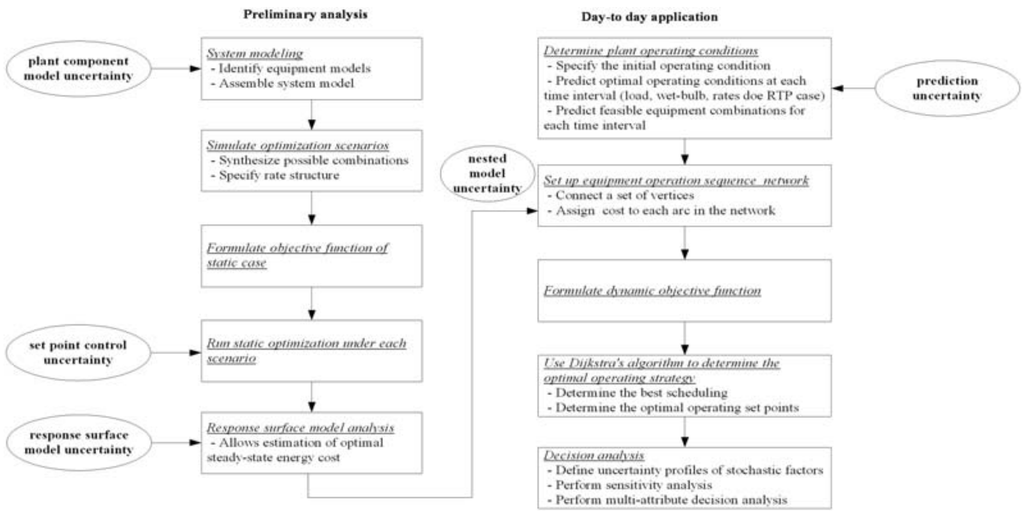

Figure 1. Framework of the proposed decision analysis methodology.

Uncertainty from Nested Regression Models

Mismatch between actual plant performance and that predicted from a model of limited complexity is one major contribution to the plant model's inherent uncertainty. The other major source of uncertainty is that incomplete data and/or erroneous data may be used to fit the model. This could arise if data points used to identify the model coefficients did not cover the entire operation region of the plant. In such a case, model predictions could be biased if used to predict plant performance outside this region.

Usually, these variations and disturbances in the model can be dealt with by including an error term in the model. A general form of a multivariable ordinary least-squares (OLS) regression model with one response variable *y* and *p* regressor variables *x* can be expressed as (Reddy and Anderson 2002):

$$y_{(n,1)} = x_{(n,p)}\beta_{(p,1)} + \varepsilon_{(n,1)}(0,\sigma^2) \tag{1}$$

where the subscripts indicate the number of rows and columns of the vectors or matrices. The random error term is assumed to have a normal distribution with variance σ^2 and no bias (mean is zero). Since the terms in the models are assumed to be independent variables, their variances can be assumed to be additive.

Energy consumption estimates at each hour are obtained from the response surface model for each system operating mode, which allows predicting the minimal operating cost for given values of the forcing variables (see Jiang and Reddy [2007]). Since the forcing variables and the models upon which the engineering optimization is based have uncertainties, the energy consumption estimate also will have an associated error value, which, under the assumption of additive errors and OLS regression models, can be modeled as the combination of the component uncertainty (σ_c) and the response surface model uncertainty (σ_r).

$$P_{ele}(P_{gas}) = f(Q_{ch},T_{wb},R_{ele},R_{gas}) + \varepsilon(0,\sigma^2) \tag{2a}$$

where

ε = error of energy consumption estimate:

$$\varepsilon(0,\sigma^2) = \varepsilon_c(0,\sigma_c^2) + \varepsilon_r(0,\sigma_r^2) \; ; \tag{2b}$$

ε_r = error of the response surface model;

ε_c = total error of the models for a specific group of components, which can be expressed by

$$\varepsilon_c(0,\sigma_c^2) = \varepsilon_{CH}(0,\sigma_{CH}^2) + \varepsilon_{CT}(0,\sigma_{CT}^2) + \dots , \tag{2c}$$

where ε_{CH} is the error of the chiller model, ε_{CT} is the error of the cooling tower model; and

$$\sigma_c^2 = \sigma_{CH}^2 + \sigma_{CT}^2 + \dots . \tag{2d}$$

Uncertainty Due to Control Deadbands

The effect of control implementation uncertainty on the optimal dynamic scheduling cost function is also to be considered, since control policies are usually implemented by lower-level feedback controllers whose setpoint tracking response will generally be imperfect due to actuator constraints and unmodeled time-varying behavior, nonlinearities (such as valve sticking), and disturbances (Ma et al. 1999). The control implementation uncertainty on the chilled-water temperature and the fan speed can be represented by upper and lower bounds. Based on the typical actuator constraints for these two variables, we would expect the accuracy of the temperature control to be in the range of 0.5°C~1.5°C and 2%~5% of revolutions per minute for accuracy of the fan speed control.

Prediction Uncertainty in the Time-Variant Forcing Functions

It is unrealistic to expect that perfectly accurate values for building thermal load profile, wet-bulb temperature, and electricity rate for RTP will be available for the entire planning horizon (assumed to be 12 hours). For example, the prediction of RTP rates can become very uncertain due to the announcement of variations in utility prices on short notice. These variations are typically broadcast the morning in which they take effect and can be triggered by higher-than-expected ambient temperatures (Henze and Krarti 1999). One would expect that the costs incurred will increase with decreasing accuracy of predictions. So the optimal strategy proposed must be robust enough to be used under these disturbances and errors.

Different forecasting methods provide different degrees of accuracy. The prediction models used range from the very simple (such as the unbiased random walk model) and the less simple (such as bin predictor and harmonic models) to the complex (such as the autoregressive neural network [Henze 1995]). A simple algorithm for forecasting building loads (either cooling or electrical) was proposed by Seem et al. (1989) and extended and validated by Seem and Braun (1991). The "average" time-of-day and time-of-week trends are modeled using a lookup table with time of day and type of day as the deterministic input variables. Entries in the table are updated using an exponentially weighted moving-average model. The stochastic portion of the forecast estimates errors using a third-order autoregressive model. The validation results show very good agreement between the predicted and measured demand. Four typical load prediction uncertainty models have been proposed by Henze and Krati (1999) to determine the effect of forecasting uncertainty on the cost savings performance of a predictive optimal controller for thermal storage systems. Only the first two load prediction uncertainty models were considered in this research since they were felt to be more applicable. These are described later in this paper.

Sensitivity Analysis

Monte Carlo simulation is a convenient approach to study the robustness of the identified optimal operating strategy and to evaluate the relative importance of different sources of uncertainty. It can be applied to optimal deterministic operating strategies in order to study the uncertainty effects due to (1) model-inherent uncertainty (response surface model error and component model error), (2) model-inherent uncertainty and load prediction uncertainty, and (3) model-inherent uncertainty, load prediction uncertainty, and control uncertainty. By comparing the system operating cost under different scenarios, the relative importance of these uncertainties, modeled as shown in Equation 2, can be investigated.

Decision Analysis with Multi-Attribute Utility Framework

Different operating strategies may have different uncertainty profiles with different robustness levels. For example, the expected value of the operating cost of Strategy 1 may be less than that of Strategy 2 but is less robust, i.e., it may exhibit more variability in the presence of various uncertainties listed above. One manifestation of this is that the optimal operating strategy determined based on the assumption of perfectly accurate hourly cooling load predictions could lead to a situation where the actual cooling load is not being met and a penalty is incurred by way of inadequate cooling capability of the cooling plant. This might cause a severe problem in instances that require strict temperature or humidity control. In such instances, the plant operator may want to avoid the risk of loss of cooling by running more chillers than those needed optimally, even if this strategy results in higher energy cost. Maximizing the savings and minimizing the exposure to risk are two different attributes; how to trade off between these objectives is basically a decision analysis problem.

The outcome of a decision problem consists of a number of specific attributes (in our case, expected value of operating cost, standard deviation of operating cost, and probability of loss of cooling capability). One is then faced with the problem of balancing all the outcomes so as to reach the "best decision." Different people have different personal objectives, and this needs to be considered. Further, manipulation of these multi-objective decision problems can be extremely complex, depending on the size of the problem and the degree of dependence among the various objectives, and is an area of active research (Vignaux 2005). A relatively simple model to treat such problems is based on the multi-attribute utility function (proposed by Churchman et al. [1957] about 50 years ago), which evaluates the diverse set of objectives in terms of weighted utility values and selects the alternative that best balances the competing objectives. A utility weight (or function) is first assigned to each attribute or outcome that reflects its importance to the individual. Subsequently, a multi-attribute utility function is defined, which combines these different weighted attributes into a model that captures the specific risk attitude of the decision maker. This method is the one adopted in this research.

Consider an outcome with a number of attributes ($x_1, \ldots, x_m$). Let us designate a utility function with respect to a given attribute as $U_i(x_i)$. The multi-attribute additive utility function approach is simply a linear model with a weight k_i assigned to each attribute such that the utility of this outcome is modeled (Clemen and Reilly 2001) as

$$U(x_1, \ldots, x_m) = k_1 U_1(x_1) + \ldots + k_m U_m(x_m) = \sum_{i=1}^{m} k_i U_i(x_i) , \qquad (3)$$

where the weights are normalized such that $k_1 + k_2 + \ldots + k_m = 1$. The utility function of each attribute also needs to be normalized such that values of 0 and 1 are assigned to the worst and best levels of that particular attribute. Note that the multi-attribute utility function approach presumes that the utility functions for different attributes are independent of each other and that there is no interaction among the attributes.

Based on the above model, the utility of the outcome can be determined for each alternative, and the one with the greatest expected operator utility value is the preferred choice. Alternatively, different choices can be compared using indifference curves, which are graphs that plot combinations of two attributes such that points on the curve would imply the decision-maker has no preference for one combination over the other (Mehrez and Gafni 1990). However, even a small change in the weights of one of the attributes makes one choice clearly preferable over the other. Since assigning weights to attributes is difficult in practice, this approach offers both convenience as well as clarity in making decisions, as we shall demonstrate later in this paper.

RESULTS OF DECISION ANALYSIS UNDER RTP

The exact same hybrid cooling plant with one absorption chiller and two vapor compression chillers as described in the companion paper (Jiang and Reddy 2007) is selected to illustrate the proposed decision analysis methodology. Further, the same six diurnal conditions (RTP-1 to RTP-6), as well as the diurnal profiles for cooling loads, wet-bulb temperature, and electricity rates, are also assumed. Recall that RTP has no demand charge and is the sum of the operating cost of equipment under steady-state operation and the cost of additional energy use during equipment start-up. The specifics of these 6 RTP cases are shown in Table 7 of the companion paper (Jiang and Reddy 2007).

Sensitivity Analysis

As described earlier, the model-inherent uncertainty for each group is characterized by a normal random error of zero bias and σ standard deviation represented by $\varepsilon(0,\sigma^2)$. The value of σ is computed following Equation 2 using the individual equipment model standard deviations of the residuals of each regressed model (these values are given in Table 2 of the companion paper [Jiang and Reddy 2007]). Table 1 summarizes the coefficient of variation of the root mean squared error (CV-RMSE) values of the regression models for the five operating groups of the hybrid cooling plant (G1–G5). The two load prediction uncertainty models, random and autocorrelated errors assumed here, are based on Henze and Krarti (1999) (see Jiang [2005] for more details).

Unbiased, Uncorrelated Gaussian Noise. Normally distributed deviations around the true value are assumed, whose variance grows linearly with the look-ahead horizon. This corresponds to a predictor whose prediction performance deteriorates linearly with time. Mathematically, this noise model can be expressed as:

$$y_{t,i} = x_{t+i}[1 + \varepsilon(0,\sigma_i^2)] \tag{4}$$

$$\sigma_i = p\frac{i}{L} \tag{5}$$

where

$y_{t,i}$ = output variable at hour t and look-ahead hourly interval i

x_t = input at hour t

L = planning horizon (in our case, L = 12 hours)

p = a scalar characterizing the noise magnitude

Different noise levels were investigated, p = {0, 0.05, 0.10, 0.15, 0.2}, though in practice one would expect $p \approx 0.05$. Figure 2 illustrates the extent to which the load prediction error keeps

Table 1. Summary of CV-RMSE Values for All Regression Models for the Hybrid Cooling Plant (Groups Indicate Feasible Configurations of Various Components Described in Jiang and Reddy [2007])

Group	Description	Response Variable	CV-RMSE (%)	
			Component Model	Response Surface Model
G1	One VC[a] chiller and one cooling tower	P_{ele}	4.0	0.14
		P_{gas}	—	—
G2	One absorption chiller and one cooling tower	P_{ele}	3.0	2.9
		P_{gas}	5.0	0.27
G3	Two VC chillers and two cooling towers	P_{ele}	5.7	1.1
		P_{gas}	—	—
G4	One VC chiller and one absorption chiller and two cooling towers	P_{ele}	5.0	4.7
		P_{gas}	5.0	1.2
G5	All three chillers and all three cooling towers	P_{ele}	5.7	4.2
		P_{gas}	5.0	2.3

a. VC = vapor compression.

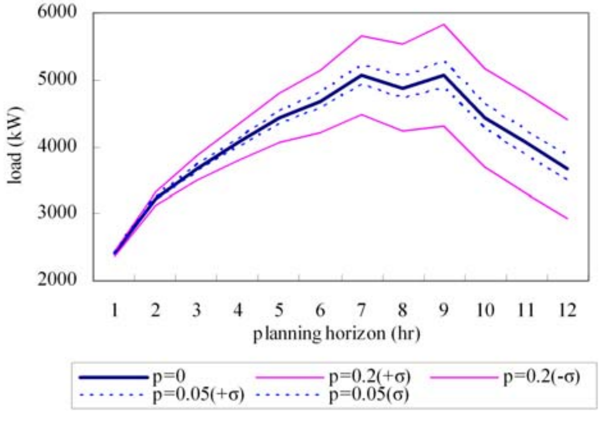

Figure 2. Load prediction versus planning horizon for different noise levels for uncorrelated Gaussian noise on a typical hot day in Philadelphia.

increasing along with the look-ahead hour. At the end of 12 hours, the error bands represented by $p = 0.2$ are 20% of the mean value. As expected, the load prediction error for $p = 0.2$ is much greater than that for $p = 0.05$.

Unbiased, Correlated Gaussian Noise. If the last prediction error affects the next prediction error, the process is said to have *correlated noise*. In other words, if one overpredicts the load for the current hour, it is likely that the load would be equally overpredicted for subsequent hours. In this research, a first-order autoregressive procedure was used:

Table 2. Sensitivity Analysis Results with ODS under Diurnal Condition RTP-6 (Based on 10^4 Monte Carlo Simulation Trials)

$(\varepsilon_m, \varepsilon_l)$		(0, 0)	$(\varepsilon_m, 0)$	$(\varepsilon_m, 0.05)$	$(\varepsilon_m, 0.10)$	$(\varepsilon_m, 0.15)$	$(\varepsilon_m, 0.2)$	$(\varepsilon_{m\cdot 2}, 0)$	$(\varepsilon_{m\cdot 2}, 0.05)$	$(\varepsilon_{m\cdot 2}, 0.2)$
Uncorrelated Gaussian Noise										
Operating cost of plant	EV ($)	2267.7	2267.2	2267.8	2268.6	2270.6	2274.2	2266.9	2267.8	2273.2
	STD ($)	0	35.4	39.1	47.4	59.3	70.8	70.7	73.3	93.2
	CV-STD (%)	0	1.6	1.7	2.1	2.6	3.1	3.1	3.2	4.1
PLC (%)		0	0	0	0.17	1.9	4.5	0	0	4.6
Correlated Gaussian Noise										
Operating cost of plant	EV ($)	2267.7	2267.2	2268.0	2268.7	2267.4	2267.1	2266.9	2268.0	2265.8
	STD ($)	0	35.4	49.4	76.3	102.6	129.8	70.7	79.3	142.5
	CV-STD (%)	0	1.6	2.2	3.4	4.5	5.7	3.1	3.5	6.3
PLC (%)		0	0	0	0.69	4.1	7.1	0	0	6.9

Note: $(\varepsilon_m, \varepsilon_l,)$ describes the uncertainty profile, ε_m is the actual model uncertainty in the hybrid cooling plant, $\varepsilon_{m\cdot 2}$ denotes the model uncertainty is two times the actual model uncertainty; ε_l is the load prediction uncertainty with values assumed to be 5%, 10%, 15%, and 20%. EV is the expected value of operating cost, SD is the standard deviation of operating cost, and PLC is the probability of loss of cooling.

$$y_{t,i} = x_{t+i}[1 + \sqrt{1 - C^2} \cdot \varepsilon(0, \sigma_i^2)] + C \cdot (y_{t,i-1} - x_{t+i-1}) \cdot \frac{\sigma_i}{\sigma_{i-1}} \tag{6}$$

where σ_i is defined by Equation 2 and C is the correlation coefficient. The greater the numerical value of this coefficient, the more strongly does the past prediction error influence the current prediction error. In this study, we have chosen to model a strong correlation with $C = 0.8$ as an upper bound. The prediction error with correlated Gaussian noise model was found to be somewhat greater than that with the uncorrelated Gaussian noise model, but not significantly so, even for $C = 0.8$ (Jiang 2005).

Latin Hypercube Monte Carlo (LHMC) simulation with 10^4 trials was then applied to the optimal deterministic operating strategies (ODS) under different diurnal conditions. LHMC sampling is a form of stratified sampling that can be applied to multiple variables. The method is commonly used to reduce the number of runs necessary for a standard Monte Carlo simulation to achieve a reasonably accurate random distribution (Vose 1996). Among all the diurnal conditions, RTP-6 has the greatest operating cost difference between static optimization and the least-cost path algorithm (Jiang 2005). Therefore, RTP-6, with the highest model-inherent uncertainty, is selected for the sensitivity analysis in order to establish an upper bound for how uncertainty impacts the deterministic results. The results are summarized in Table 2 for both uncorrelated and correlated Gaussian noise, from which the following observations can be made.

- Expected values of the operating cost of the hybrid cooling plant can be taken to be statistically constant under all the uncertainty profiles since the LHMC runs assumed zero bias in the error term (one can note small differences, less than 0.1%, that are due to the random number generation process).
- As expected, the coefficient of variation of the standard deviation (CV-STD) values of operating cost increase with increasing model error. When the model error is doubled, the CV-STD values are approximately doubled as well. Also, as expected, CV-STD values of the operating cost increase with increasing load prediction uncertainty (Figures 3 and 4).
- CV-STDs are greater under the assumption of correlated Gaussian noise than those under uncorrelated Gaussian noise (see Figures 3 and 4). The difference increases significantly with increasing load prediction uncertainty but remains almost constant with increasing model uncertainty.

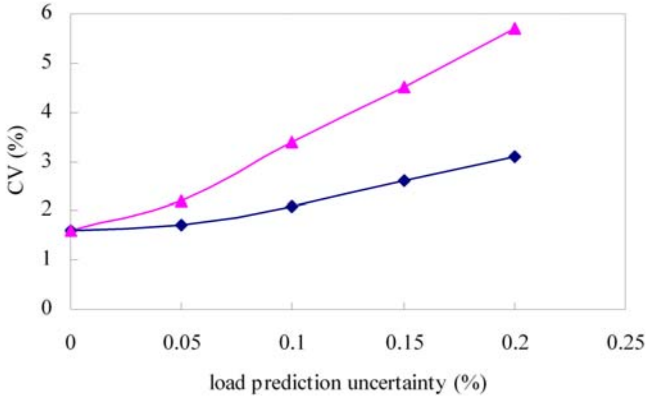

Figure 3. CV-STD vs. load prediction uncertainty for RTP-6 with model uncertainty ε_m.

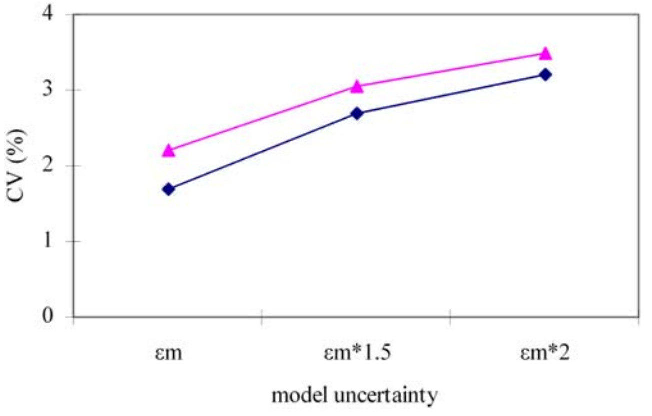

Figure 4. CV-STD vs. model uncertainty for RTP-6 with load prediction uncertainty of 5%.

- Regarding the CV-STD values of the operating cost, model-inherent uncertainty is relatively more important than the load prediction uncertainty. For example, if the uncertainty profile is changed from (ε_m, 0.05) to ($\varepsilon_m \cdot 2$, 0.05) (ε_m is approximately 0.05 under RTP-6), the CV-STD of the minimal plant operating cost is increased from 1.7% to 3.2% (Table 2). On the other hand, if the load prediction error is doubled while the model uncertainty is kept constant, the CV-STD of operating cost is only increased to 2.1%. However, the model uncertainty has less effect on probability of loss of cooling capability than the load prediction error.

In summary, under a practical uncertainty condition of (ε_m, 0.05), model-inherent uncertainty and load prediction uncertainty seem to have little effect on the overall operating cost of the hybrid cooling plant under optimal operating strategy with the CV-STD being around 2%.

Besides the model-inherent uncertainty and load prediction uncertainty, we also studied the effect of the control uncertainty due to the control dead band. The sensitivity in the cost of energy consumption of the hybrid cooling plant to the two control variables when only two centrifugal chillers are in operation can be gleaned from Figures 5 and 6. The two variables are loading fraction (controlled by the chilled water supply temperature) and condenser inlet water temperature (controlled by the cooling tower fan speed). From Figures 5 and 6, we can state that the minimum cost is fairly insensitive to the loading fraction and the wet-bulb temperature T_{wb} since the curves are fairly flat (a change of only ±\$1.50 over \$30, i.e., a ±5% change over a large chiller loading fraction range of 0.2–0.8). Therefore, we conclude that the uncertainty due to control deadbands does not have much influence on the optimal strategy. Hence, the effect of control deadbands need not be explicitly considered in the decision analysis study that follows.

Decision Analysis with Multi-Attribute Utility Framework

The above sensitivity analysis revealed that different uncertainty scenarios have different expected values (EV) and CV-STD values for the cost of operating the hybrid cooling plant. Also discussed previously is the issue of loss of cooling capability that could occur due to uncertainty in the building cooling load prediction. This example assumed that the maximum power (or gas) consumption of each chiller is the same as the chiller power (or gas) consumption under ARI rated conditions (i.e., $P_{ele,ch,k,max} = P_{ele\_rated,ch,k}$) (ARI 2000, 2003). The probability of loss of cooling capability is determined under each uncertainty scenario following the procedure described below.

- Find the maximum cooling capacity of each chiller under specified weather condition following the procedure shown in Figure 7. An iteration loop that includes an energy balance of the cooling plant, chiller model, and cooling tower model is used to calculate the maximum cooling capacity of each chiller. The sum of the maximum cooling capacity of each chiller ($Q_{ch,k,max}$) yields the maximum cooling capacity of the whole hybrid cooling plant at a given hour:

$$Q_{ch,max} = \sum_{k=1}^{3} Q_{ch,k\ max}$$

  (for full details see Jiang [2005]).
- Under each specified uncertainty scenario and diurnal condition, LHMC simulation is used to generate numerous trials of the building cooling load profile. Subsequently, for each trial, one can again generate numerous cases using LHMC to simulate the effect of component model errors on the energy consumption of the chillers under ODS as determined by the least-cost

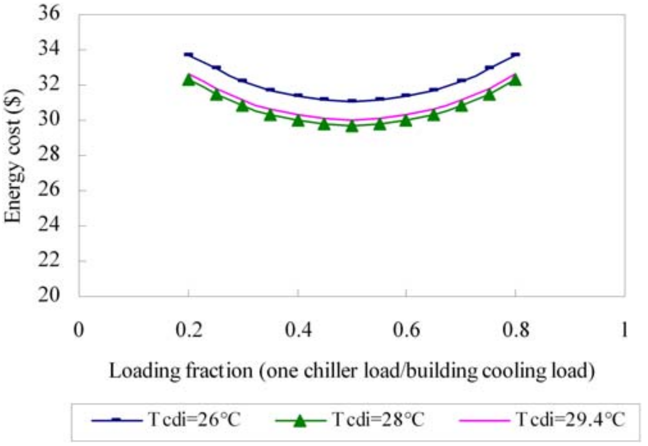

Figure 5. Energy cost of the hybrid cooling plant vs. loading fraction at different T_{cdi} when two centrifugal chillers are in operation at T_{wb} = 24°C, building cooling load = 3165 kW, R_{ele} = \$0.05 kW/h, and R_{gas} = \$0.4/therm.

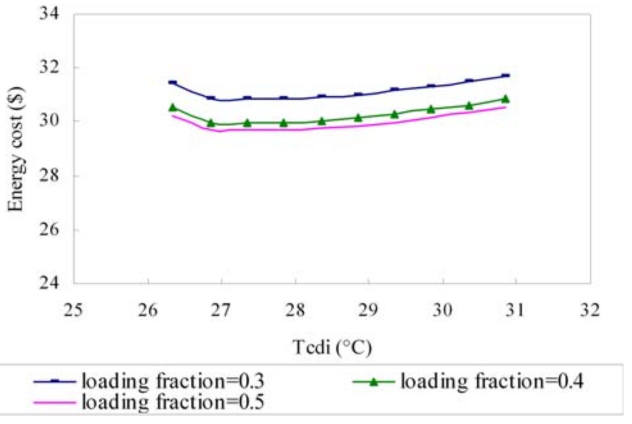

Figure 6. Energy cost of the hybrid cooling plant vs. T_{cdi} at different loading fraction with two chillers equally loaded when two centrifugal chillers are in operation at T_{wb} = 24°C, building cooling load = 3165 kW, R_{ele} = \$0.05/kWh, and R_{gas} = \$0.4/therm.

path algorithm. This nested uncertainty computation thus captures the effect of both sources of uncertainty, load prediction and component model.

- Only if $Q_{ch,bldg,\max} > Q_{ch,\max}$ and $P_{ele,ch} > P_{ele,ch,\max}$ does a situation of loss of cooling capability occur. Note that $P_{ele,ch,\max}$ is defined as

$$P_{ele,ch,\max} = \sum_{k=1}^{3} P_{ele,ch,k\,\max}\,.$$

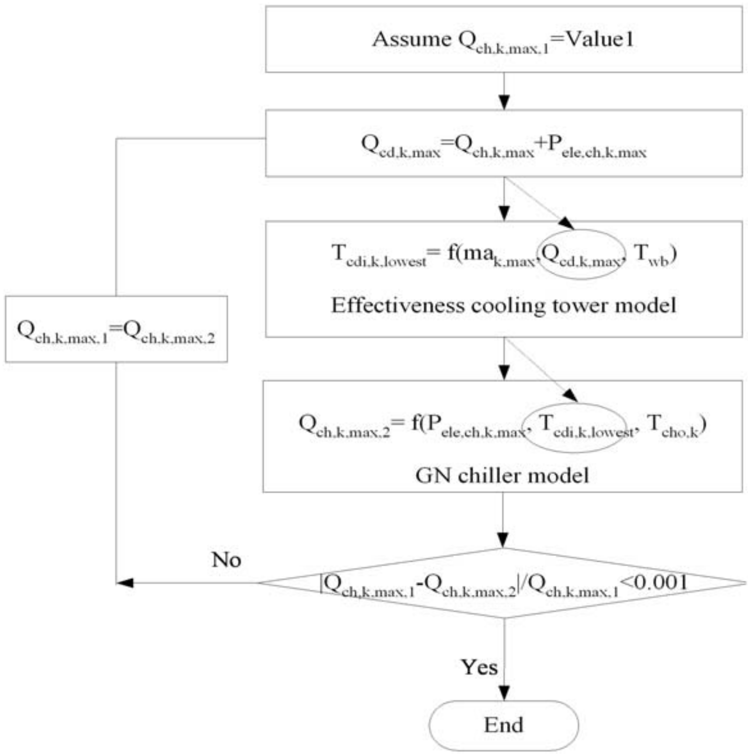

Figure 7. Procedure for finding maximum chiller cooling capacity.

The total number of occurrences of loss of cooling capability is divided by the total number of simulation trials to yield the probability of the loss of cooling capability (PLC) under specified uncertainty scenario and optimal operating strategy. PLC values for diurnal condition RTP-6 are also shown in Table 2.

Risk-Averse Strategy

Certain circumstances warrant operating the cooling plant differently than as suggested by the optimal deterministic strategy (ODS). For example, the activities inside the building may demand that the required cooling load always be met (such as in a pharmaceutical company, for example). Hence, one would like to have excess cooling capability at all times even if this results in extra operating cost. One could identify different types of risk-averse strategies (RAS). For example, one could simply run one additional chiller than the number suggested by ODS at all times. Alternatively, one could bring an extra chiller online only when the reserve capacity is less than a certain amount, say 10%. In this study, we selected the former scenario since this paper is meant to be conceptual in nature. Under RTP-6, this means that all three chillers keep running during the whole 12-hour planning horizon.

Sensitivity analysis is repeated for RAS under RTP-6, and results are summarized in Table 3. We note that there is no loss of cooling under any of the uncertainty scenarios, but the downside is that operating cost is increased. In this decision analysis problem, three attributes are considered: EV, CV-STD of operating cost, and PLC, referred to as x_1, x_2, and x_3, respectively, in Equation 3. The CV-STD is a measure of uncertainty surrounding the EV. In a decision frame-

Table 3. Sensitivity Analysis Results with RAS under Diurnal Condition RTP-6 (Based on 10^4 Monte Carlo Simulation Trials)

$(\varepsilon_m,\varepsilon_l)$		(0,0)	$(\varepsilon_m,0)$	$(\varepsilon_m,0.05)$	$(\varepsilon_m,0.10)$	$(\varepsilon_m,0.15)$	$(\varepsilon_m,0.2)$	$(\varepsilon_{m*2},0)$	$(\varepsilon_{m*2},0.05)$	$(\varepsilon_{m*2},0.2)$
Uncorrelated Gaussian Noise										
Operating cost of plant	EV ($)	2919.1	2921.7	2922.0	2922.5	2923.2	2924.2	2921.6	2922.0	2924.2
	STD ($)	0	37.6	39.6	45.3	53.4	62.9	75.2	76.2	90.6
	CV-STD (%)	0	1.3	1.4	1.6	1.8	2.2	2.6	2.6	3.1
PLC (%)		0	0	0	0	0	0	0	0	0
Correlated Gaussian Noise										
Operating cost of plant	EV ($)	2919.1	2921.7	2922.1	2922.7	2923.5	2924.6	2921.6	2922.1	2924.6
	STD ($)	0	37.6	46.0	65.0	87.8	111.9	75.2	79.7	129.5
	CV-STD (%)	0	1.3	1.6	2.2	3.0	3.8	2.6	2.7	4.4
PLC (%)		0	0	0	0	0	0	0	0	0

work, higher uncertainty results in higher risk, and so this attribute needs to also figure in the utility function. Since these three attributes do not interact with affect each other, it is reasonable to assume them to be additive independent, which allows for the use of an additive multi-attribute utility function of the forms:

$$U_{ODS}(EV, CV\text{-}STD, PLC) = k_{1,ODS}U_{1,ODS}(EV) + k_{2,ODS}U_{2,ODS}(CV\text{-}STD) + k_{3,ODS}U_{3,ODS}(PLC) \quad (7)$$

$$U_{RAS}(EV, CV\text{-}STD, PLC) = k_{1,RAS}U_{1,RAS}(EV) + k_{2,RAS}U_{2,RAS}(CV\text{-}STD) + k_{3,RAS}U_{3,RAS}(PLC) \quad (8)$$

where EV is the expected value of total operating cost over the planning horizon, CV-STD is the coefficient of variation of the standard deviation of the total operating cost from all the Monte Carlo trials, PLC is the probability of loss of cooling capability, and k denotes the three weights normalized such that $k_{1,ODS} + k_{2,ODS} + k_{3,ODS}$ and $k_{1,RAS} + k_{2,RAS} + k_{3,RAS} = 1$.

The benefit of using the additive utility function is that it allows us to easily assess the impact of the individual utility and weights (Clemen and Reilly 2001). Here we assume that all the attributes are risk neutral (risk neutrality is characterized by a linear utility function). The terms U_1, U_2, and U_3 under both ODS and RAS are normalized by dividing them by the differences of worst value and best value of the three attributes, respectively, so that these values are bounded by values of 0 and 1. The normalized utilities of the three attributes at other levels, ranging from worst to best, can be determined by linear interpolation. Table 4 lists some utility values of the three attributes for RTP-6 under five different combinations or cases of model and load prediction uncertainties. Hence, knowing the weight k, one can calculate the overall utility under different uncertainty scenarios. However, it is difficult to assign specific values of the weights

Table 4. Utility Values for ODS and RAS Strategies under Uncorrelated Gaussian Noise Assumption and Diurnal Condition RTP-6 Normalized Such that They are between 0 and 1, Representing the Worst and Best Values, Respectively; Cases Are Defined in Table 1

		Case 1	**Case 2**	**Case 3**	**Case 4**	**Case 5**
	(ε_m,ε_l)	**(ε_m,0)**	**(ε_m,0.05)**	**(ε_m,0.10)**	**(ε_m,0.15)**	**(ε_m,0.2)**
	U(EV)	1.00	1.00	1.00	1.00	1.00
ODS	U(CV-STD)	0.88	0.84	0.68	0.48	0.28
	U(PLC)	1.00	1.00	0.96	0.58	0.00
	U(EV)	0.00	0.00	0.00	0.00	0.00
RAS	U(CV-STD)	1.00	0.88	0.64	0.32	0.00
	U(PLC)	1.00	1.00	1.00	1.00	1.00

since they are application-specific. If occasional loss of cooling can be tolerated, then k_3 can be either set to zero or given very low weight. Also, a general heuristic guideline is that the weight for the expected value (k_1) is usually two to four times greater than that of the variability of the predicted operating cost (k_2) (Clemen and Reilly 2001).

The values for the multi-attribute utility functions $U_{ODS}(EV, CV\text{–}STD, PLC)$ and $U_{RAS}(EV, CV\text{–}STD, PLC$ can be calculated based on the results of Tables 2 and 3. The model uncertainty is found to have little effect on the probability of loss of cooling capability and is not considered in the following decision analysis; fixed realistic values of ε_m shown in Table 1 are assumed. All points on the indifference curves have the same utility and, hence, separate the regions of preference between ODS and RAS. These are shown in Figure 8 for the five cases defined in Table 4. These plots are easily generated by equating the right-hand terms of Equations 7 and 8 and inserting the appropriate values from Table 4. Thus, for example, if k_2 is taken to be 0.2 (a typical value), we can immediately conclude from the figure that ODS is preferable to RAS provided $k_1 > 0.375$ (which, in turn, implies $k_3 = 1 - 0.2 - 0.375 = 0.425$). The operator is typically likely to consider the attribute EV to have an importance level higher than this value. The threshold values provide a convenient means of determining whether one strategy is clearly preferred over the other or whether precise estimates of the attribute weights are required to select an operating strategy.

Figure 8 indicates that for cases 1 to 3 (load prediction error of 0%–10%), ODS is clearly preferable and that it would, most likely, be so even for cases 4 and 5 (load prediction errors of 15% and 20%). From Figure 9, which was generated assuming $k_2 = 0$ (i.e., no weight being given to variability of the operating cost), it is clear that the utility curve is steeper as k_1 decreases. This means that the load prediction error has a more profound effect on the utility function when the expected value of operating cost has a lower weight; this is consistent with our earlier observation that the load prediction uncertainty affects the loss of cooling capability. We note that RAS is only preferred under a limited set of conditions. While the exact location in parameter space of the switchover point between the two strategies may change from application to application, this approach supports the idea that well-characterized systems may be operated under ODS except when the load prediction uncertainty is higher than 15%. Further research examining a broader spectrum of system configurations would be necessary to confirm this.

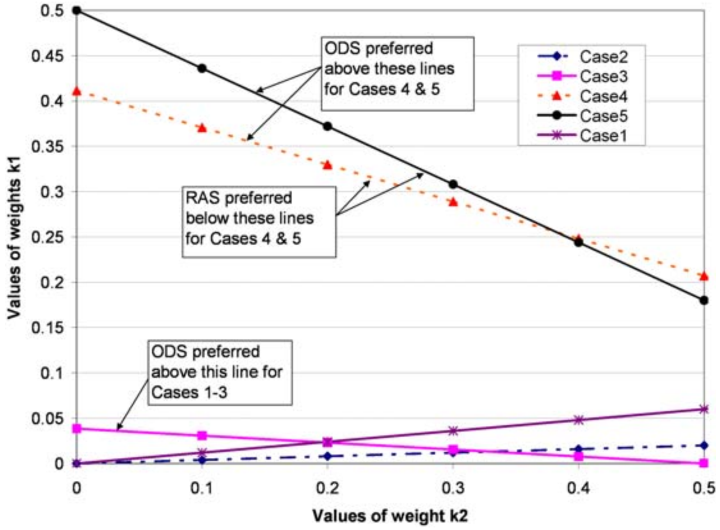

Figure 8. Indifference plots for the five cases shown in Table 4 corresponding to uncorrelated Gaussian noise and RTP-6. These indicate the preferred operating strategies for different combinations of attribute weights k_1 and k_2. Thus, for case 5 and for a selected value of k_2 = 0.2, the ODS is superior if the operator deems $U(EV)$ to have a weight k_1 > 0.375.

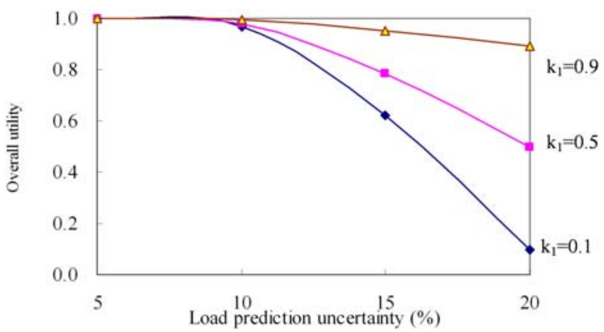

Figure 9. Overall utility of ODS vs. load prediction uncertainty for RTP-6 assuming k_2 = 0.

Finally, we noticed that the results of the decision analysis are very much dependent on the relative cooling load capacity and building load profile. For verification, the decision analysis was repeated for RTP-3, which has a lower building cooling load profile than RTP-6 (Jiang 2005). The sensitivity analysis results revealed that ODS is always able to provide sufficient cooling unless the load prediction uncertainty is greater than 50%, which is unrealistically high. Therefore, the probability of loss of cooling capability does not need to be considered under condition RTP-3, and ODS is always the preferred choice. Such conclusions are of course specific to the system configurations and scenarios selected. The above presentation serves to illustrate how circumstance-specific operating guidelines can be generated following the decision analysis approach advocated in this paper.

DECISION ANALYSIS UNDER TOU WITH ELECTRIC DEMAND

Sensitivity Analysis

TOU rate structure consists of two time periods—on-peak and off-peak hours, both of which have different electricity rate profiles. Consistent with the assumptions made in the companion paper (Jiang and Reddy 2007), we assume the planning horizon (from 9:00 to 20:00) to be on-peak hours with assumed electricity rates of \$0.1/kWh and \$0.15/kWh and demand rates of \$5/kW and \$20/kW. The gas rates assumed are \$0.4/therm and \$1.2/therm, representative of low and high values. Recall that the TOU investigation involved 16 conditions (summarized in Table 9 of the companion paper), which include various combinations of hot and mild days and the two electricity rates, electricity demand rates, and gas rates stated here. In this paper, we shall limit the presentation of the sensitivity analysis to TOU-2 (hot day load profile, \$0.1/kWh for electricity use, \$5/kW for electric demand, and \$1.2/therm for gas) and TOU-7 (hot day, \$0.15/kWh, \$20/kW, and \$0.4/therm), which among all the scenarios under high load profiles have the most different scheduling strategies since the relative electricity to gas costs are so different.

The effect of the same sources of uncertainty (namely, model uncertainty and load prediction uncertainty) is also studied under a TOU electric rate structure. Further, an additional quantity, called *load variability*, is introduced to characterize the load variation over a month due to deterministic effects, such as outdoor temperature variation during a given month. *TOU with demand charge* usually refers to monthly energy cost. The load profiles used in our study are typical daily load profiles. It is possible to run the optimization algorithm and perform subsequent uncertainty analysis over an entire month using the monthly load profiles of each day generated from monthly temperature profiles for the location, obtained from, say, TMY2 weather data tapes. However, such an approach is computationally demanding. This is the reason we introduce the concept of load variability to account for deterministic load variations and operating costs of the plant over a month. A lower range in load variability of 5% is assumed for the typical hot day load profile during the summer season, where all days are uniformly hot, and a higher range value of 20% as representative of the typical mild day load profile during swing seasons. LHMC simulations (50 trials each for model uncertainty, load prediction uncertainty, and load variability result in a total of 125,000 nested trials) are performed assuming the optimal operating strategies under different diurnal conditions.

The results for TOU-2 are listed in Tables 5 and 6, while salient following observations are presented below (see Jiang [2005] for details).

- Unlike RTP, expected values of the operating cost of the hybrid cooling plant under TOU with demand charge are not constant under all the uncertainty profiles. The expected values increase with increasing uncertainty, indicating that bias is higher when the uncertainty is higher. This is due to the demand cost component in the overall operating cost.
- As expected, the CV-STD values of operating cost increase with increasing model uncertainty and increasing load prediction errors. For most cases, standard deviations (as well as CV-STDs of the operating costs) are greater under the assumption of correlated Gaussian noise than those under uncorrelated Gaussian noise.
- Model uncertainty was found to have little effect on PLC. For example, the probability of loss of cooling capability remains constant even when the model uncertainty is changed from 0 to ε_m or from ε_m to $\varepsilon_{m \cdot 2}$. Load prediction uncertainty does affect the PLC values but less so than the load variability factor. The load variability level was found to have the most effect on the CV-STD values of the operating cost and PLC values.

Table 5. Summary of Sensitivity Analysis Results with ODS under Uncorrelated Gaussian Noise Assumption and Diurnal Condition TOU-2 (Based on 125,000 Monte Carlo Simulation Runs)

Load Variability		**No Variability in Load (0%)**	**Low Variability in Load (5%)**			
(ε_m, ε_l)		**(0, 0)**	**(0, 0)**	**(ε_m, 0)**	**(ε_m, 0.10)**	**($\varepsilon_{m\cdot 2}$, 0.10)**
Operating cost of plant	EV ($)	1685.1	1754.7	1903.5	1978.2	2223.6
	SD ($)	0	89.2	172.4	228.1	335.1
	CV(%)	0	5.1	9.1	11.5	15.1
PLC (%)		0	0	0	0.4	0.4

Table 6. Summary of Sensitivity Analysis Results with RAS under Uncorrelated Gaussian Noise Assumption and Diurnal Condition TOU-2 (Based on 125,000 Monte Carlo Simulation Runs)

Load Variability		**No Variability in Load (0%)**	**Low Variability in Load (5%)**		
(ε_m, ε_l)		**(0, 0)**	**(0, 0)**	**(ε_m, 0)**	**(ε_m, 0.10)**
Operating cost of plant	EV ($)	1991.7	2076.8	2232.7	2283.8
	SD ($)	0	78.7	182.1	213.7
	CV(%)	0	3.8	8.2	9.4
PLC (%)		0	0	0	0

Decision Analysis with Multi-Attribute Utility Framework

The RAS under TOU is also defined as operating one more chiller than the number indicated by ODS. Under both TOU-2 and TOU-7, this means that all three chillers are kept running over the whole 12-hour planning horizon. Sensitivity analysis was done for RAS for both cases, and Table 6 gives the results for TOU-2. Again, we assume that all the attributes are risk-neutral, with worst and best levels of each attribute among the ODS and RAS being bounded by values of 0 and 1. Because the model uncertainty has little effect on the probability of loss of cooling capability, it is fixed at a realistic value ε_m (see Table 1). The utility values are listed in Table 7 for two cases of load prediction uncertainty (0% and 10%) for ODS and RAS. The sensitivity analysis results allow the values for the multi-attribute utility function $U_{ODS}(EV, CV\text{–}STD, PLC)$ and $U_{RAS}(EV, CV\text{–}STD, PLC$ with assigned weight k to be calculated following Equations 7 and 8. Figure 10 shows the indifference curves for these two cases as was done for the RTP scenario. Even here, ODS seems to be clearly preferable in most instances. For example, an upper limit for weight k_2 would be 0.3, in which case the threshold value would be $k_1 = 0.23$ for the more critical case 2. It is unlikely that the PLC weight $k_3 > (1 - 0.3 - 0.23)$, i.e., 0.47. Hence, for this specific plant configuration and cases studied, the operator can conclude that ODS is the operating strategy of choice.

SUMMARY

Studies to date dealing with model-based optimal operation of primary building energy systems have tended to focus primarily on the aspect of engineering optimization while overlooking

Table 7. Utility Values for ODS and RAS Strategies under Uncorrelated Gaussian Noise Assumption, 5% Load Variability and Diurnal Condition TOU-2

$(\varepsilon_m, \varepsilon_l)$		**Case 1** $(\varepsilon_m, 0)$	**Case 2** $(\varepsilon_m, 0.10)$
ODS	U(EV)	1.00	0.90
	U(CV-STD)	0.90	0.63
	U(PLC)	1.00	0.90
RAS	U(EV)	0.50	0.40
	U(CV-STD)	1.00	0.86
	U(PLC)	1.00	1.00

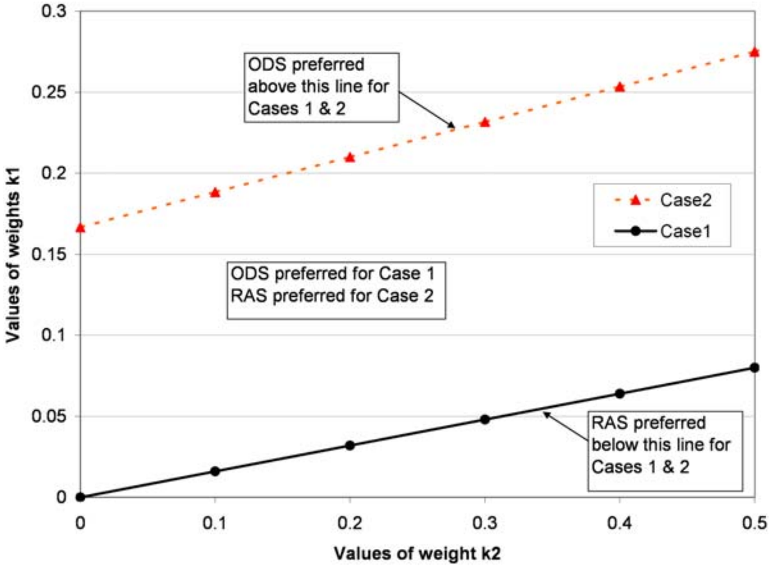

Figure 10. Same as Figure 8 but for uncorrelated Gaussian noise, 5% load variability, and TOU-2. Cases are defined in Table 7.

the equally important aspect of how to include various sources of uncertainty and individual operator risk preferences into the decision-making process. We feel that this is one of the reasons why most complex HVAC&R plants are still scheduled by humans in a heuristic manner without the aid of computer supporting tools. This paper proposed a framework that allows a systematic approach to sensitivity analysis of the various sources of uncertainty present in the dynamic scheduling of complex primary HVAC&R systems while considering individual risk preferences of the plant operator. The paper introduced the concept of evaluating non-optimal but risk-averse strategies in terms of specific attributes and proposed a multi-attribute utility function approach for decision making. Specifically, the methodology proposed involved the following steps: (1) define uncertainty profiles for different uncertainty sources; (2) perform sensitivity analysis on the ODS strategy using Monte Carlo simulations so as to compute the attributes, which in this study have been assumed to be the expected values of the operating costs (EV), the variability of operating costs (characterized by the CV-STD), and loss of cooling probability (PLC) under different uncertainty levels; (3) identify alternative risk-averse operating strategies (RAS) that may be non-optimal in terms of operating cost but are likely to reduce the risk of loss of cooling capability; (4) compute, again using Monte Carlo simulations, the effect on the same set of attributes; and (5) adopt an additive multi-attribute utility function formulation and evaluate the trade-off between ODS and RAS so that the one with higher utility value can be selected. The proposed approach is illustrated for the simple but realistic case of a hybrid cooling plant with two vapor compression chillers and one absorption chiller operated under RTP and TOU pricing signals (the very same plant assumed in the companion paper [Jiang and Reddy 2007]).

In this research, we have assumed a simple RAS scheme involving operating an additional chiller than that suggested by ODS. It is obvious that the methodology proposed is equally appropriate for more sophisticated and realistic RAS strategies that a plant operator may wish to evaluate. Further, the suggested framework would conveniently allow utility attributes other than the three selected here to be considered. The computational efficiency of the deterministic dynamic scheduling algorithm allows several thousand Monte Carlo runs to be performed very quickly. It is felt that the proposed methodology, coupling the deterministic engineering optimization within a decision analysis framework, would be more appealing to operators of complex plants, such as HVAC&R systems; building cooling, heating, and power; and distributed energy plants and would result in spurring the use of decision support tools in these application areas.

ACKNOWLEDGMENTS

We thank Dr. Itzhak Maor for providing us with detailed information of the actual hybrid cooling plant that served as the basis of the case study in this research and also for useful advice and practical insights throughout this work. Insightful comments from Drs. J. Wen, S.V. Smith, and J.R. Weggel are also acknowledged. The paper also benefited greatly from the detailed comments of the anonymous reviewers.

NOMEMCLATURE

C = correlation coefficient
CV-RMSE = coefficient of variation of the root mean square error
CV-STD = coefficient of variation of the standard deviation
DCS = dynamic chiller sequencing
EV = expected value of total operating cost over planning horizon
k = weight of utility function, equipment or component index
L = planning horizon
LHMC = Latin Hypercube Monte Carlo
ODS = cost-optimal deterministic strategy of operating the plant
OLS = ordinary least squares
P = electric or gas energy consumption per unit time interval, kWh/h or therm/h
p = scalar for the noise magnitude.
PLC = probability of loss of cooling capability
$Q_{ch,bldg}$ = building cooling load, kW
Q_{ch} = chiller cooling load, kW
RAS = risk-averse strategy of operating the plant
R_{ele} = unit cost of electricity, $/kWh
R_{gas} = unit cost of gas, $/therm
STD = standard deviation
T_{wb} = ambient wet-bulb temperature, °C
x_m = attribute
x_t = input at hour t
$U_m(x_m)$ = utility functions for attribute x_m
$y_{t,i}$ = output variable at hour t and look-ahead hourly interval i
ε = error of energy consumption estimate
ε_c = total error of the models for a specific group of components
ε_r = error of the response model
σ = standard deviation

Subscripts

c	= component	gas	= gas
CH	= chiller	max	= maximum
CT	= cooling tower	r	= response surface
ele	= electricity		

REFERENCES

ARI. 2000. *Standard 560-2000*, *Absorption Water Chilling and Water Heating Packages*. Arlington, VA: Air Conditioning and Refrigeration Institute.

ARI. 2003. *Standard 550/590-2003*, *Performance Rating of Water-Chilling Packages Using the Vapor Compression Cycle*. Arlington, VA: Air Conditioning and Refrigeration Institute.

Charnes, A., and W.W. Cooper. 1959. Chance constrained programming. *Management Science* 6(1)[October]:73–79.

Churchman, C., E. Arnoff, and R. Ackoff. 1957. *Introduction to Operations Research*. New York: John Wiley & Sons.

Clemen, R.T., and T. Reilly. 2001. *Making Hard Decisions with Decision Tools*. Pacific Grove, CA: Duxbury/Thomson Learning.

Dantzig, G.B. 1955. Linear programming under uncertainty. *Management Science* 1:197–206.

Haimes, Y.Y., and W.A. Hall. 1977. Sensitivity, responsivity, stability, and irreversibility as multiple objectives in civil systems. *Advances in Water Resources* 1(2)[December]:71–81.

Henze, G.P. 1995. Evaluation of optimal control for ice storage systems. PhD dissertation, University of Colorado, Boulder, CO.

Henze, G.P., and M. Krarti. 1999. The impact of forecasting uncertainty on the performance of a predictive optimal controller for thermal energy storage systems. *ASHRAE Transactions* 105(2):553–61.

Henze, G.P., M. Krarti, and M.J. Brandemuehl. 2003. Guidelines for improved performance of ice storage systems. *Energy and Buildings* 35(2):111–27.

Hillier, F.S., and G.J. Lieberman. 2001. *Introduction to Operations Research*, 7th ed. New York: McGraw-Hill.

Jiang, W. 2005. Framework combining static optimization, dynamic scheduling and decision analysis applicable to complex primary HVAC&R plants. PhD dissertation, Drexel University, Philadelphia, PA.

Jiang, W., and T.A. Reddy. 2007. General methodology combining engineering optimization of primary HVAC&R plants with decision analysis methods—Part I: Deterministic analysis. *HVAC&R Research* 13(1).

Ma, D.L., S.H. Chung, and R.D. Braatz. 1999. Worst-case performance analysis of optimal batch control trajectories. *AIChE Journal* 45(7):1469–76.

Mehrez, A., and A. Gafni. 1990. Evaluating health related quality of life: An indifference curve interpretation for the time trade-off technique. *Social Science & Medicine* 31(11):1281–83.

Mulvey, J.M., R.J. Vanderbei, and S.A. Zenios. 1995. Robust optimization of large-scale systems. *Operations Research* 43(2)[March-April]:264–281.

Olson, R.T. 1987. Optimal allocation of building cooling loads to chilled water equipment. PhD dissertation, University of Illinois at Urbana-Champaign, Urbana, IL.

Pistikopoulos, E.N. 1995. Uncertainty in process design and operations. *Computers and Chemical Engineering* 19, Suppl., pp. S553–63.

Reddy, T.A., and K.K. Andersen. 2002. An evaluation of classical steady-state off-line linear parameter estimation methods applied to chiller performance data. *HVAC&R Research* 8(1):101–24.

Seem, J., P. Armstrong, and C. Hancock. 1989. Comparison of seven methods for forecasting the time to return from night setback. *ASHRAE Transactions* 95(2):439–46.

Seem, J.E., and J.E. Braun. 1991. Adaptive methods for real-time forecasting of building electrical demand. *ASHRAE Transactions* 97(1):710–21.

Seferlis, P., and A.N. Hrymak. 1996. Sensitivity analysis for chemical process optimization. *Computers Chem. Eng.* 20(10):1177–1200.

Tintner, G. 1955. Stochastic linear programming with applications to agricultural economics. In H.A. Antosiewicz (ed.), *Proceedings of the Second Symposium on Linear Programming* (pp. 197–228), National Bureau of Standards, Washington, DC.

Vignaux, G.A. 2005. Multi-attribute decision problems. www.mcs.vuw.ac.nz/courses/OPRE251/2005T1/Lecture-Notes/multi.ps. Accessed April 12, 2004.

Vose, D. 1996. *Quantitative Risk Analysis: A Guide to Monte Carlo Simulation Modeling*. Chichester, New York: John Wiley and Sons.