CH-89-6-6

Applications of Optimal Control to Chilled Water Systems without Storage

J.E. Braun S.A. Klein J.W. Mitchell W.A. Beckman

ABSTRACT

In this paper, optimization techniques are applied to analyzing the optimal performance of chilled water systems that do not have significant thermal storage. The important uncontrolled variables that affect system performance and optimal control settings are identified. Control guidelines useful to plant engineers for improved control practices are developed. In addition, results and conclusions concerning both control and design under optimal control of chilled water systems are presented.

INTRODUCTION

A central cooling plant has many operating variables that may be controlled to minimize operational costs. Global optimum plant control has been studied by Marcev (1980), Arnold (1984), Sud (1984), Hackner (1984, 1985), Lau (1985), Johnson (1985), and Nugent (1988). These studies primarily demonstrated the potential savings associated with the use of optimal control. They did not produce general algorithms suitable for on-line optimal plant control nor general guidelines for near-optimal control.

In a companion paper, Braun (1989) presents two methodologies for determining the optimal control of chilled water systems without thermal storage. A component-based nonlinear optimization algorithm was developed as a simulation tool for investigating optimal system performance. A system-based methodology was also presented for near-optimal control that depends upon simplifications that reduce the complexity of the control problem.

In this paper, the component-based optimization methodology is utilized as a tool for investigating chilled water systems under optimal control. With this tool, general guidelines for near-optimal performance are developed. These guidelines are incorporated in the system-based near-optimal control methodology, but they are also important to plant engineers for improved control practices. The component-based optimization is also used in this study to compare optimal control with alternative control practices and different system configurations under optimal control. The simulation models of the equipment utilized in this study are described by Braun (1987, 1988).

SYSTEM DESCRIPTION

Figure 1 shows a general schematic of the variable-air-volume (VAV) chilled water system considered in this study. The central cooling facility, which consists of multiple centrifugal chillers, cooling towers, and pumps, provides chilled water to a number of air-handling units in order to condition air that is supplied to building zones. All energy-consuming components in the system are assumed to be electrically driven.

At any given time, it is possible to meet the cooling needs with a number of different modes of operation and setpoints. Optimal control of a system involves minimizing the total power consumption of the chillers, cooling tower fans, condenser water pumps, chilled water pumps, and the air-handling fans at each instant of time with respect to the independent continuous and discrete control variables. Discrete control variables are not continuously adjustable, but have discrete settings. The discrete control variables considered in this paper include the number of operating chillers, cooling tower cells, condenser water pumps, and chilled water pumps and the relative speeds for multi-speed fans or pumps. The independent continuous control variables considered include the chilled water and supply air set temperatures, relative water flow rates to multiple chillers (both evaporators and condensers) and multiple cooling tower cells, and the speeds for variable-speed cooling tower fans and chilled and condenser water pumps.

In addition to the independent optimization control variables, there are also local loop (dependent) controls associated with the chillers, air handlers, and chilled water pumps. All local loop controls are assumed to be ideal, such that their dynamics are not considered. Each chiller is considered to be controlled such that the specified

J.E. Braun, Senior Research Engineer, Johnson Controls Inc., Milwaukee, WI; **S.A. Klein, J.W. Mitchell,** and **W.A. Beckman,** Professors of Mechanical Engineering, University of Wisconsin-Madison.

ASHRAE Transactions 1989, V. 95, Pt. 1

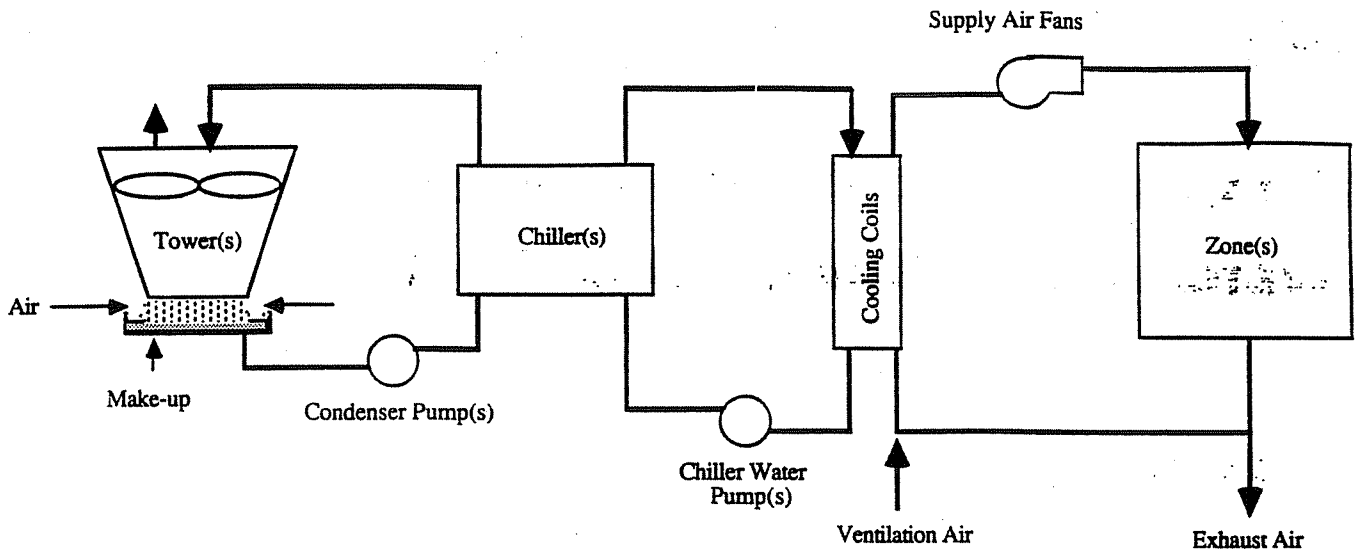

Figure 1 *Schematic of a typical chilled water system*

chilled water set temperature is maintained. The air handler local loop control involves control of both the coil water flow and fan air flow in order to maintain the prescribed supply air setpoint and zone temperature. The total requirement for the chilled water flow to the air handlers is dictated by the chilled water and supply air setpoints and the load. Control of the chilled water pumps is implemented through a local loop control that maintains a constant pressure difference between the main supply and return pipes for the air handlers. The setpoint for this pressure difference is chosen to ensure adequate distribution of flow to all air handlers and is not considered an optimization variable.

The optimal control variables change through time in response to uncontrolled variables. The uncontrolled variables are measurable quantities that may not be controlled but that affect the component outputs and/or costs, such as the load and ambient dry- and wet-bulb temperature.

The results presented in this study are primarily representative of the Dallas/Fort Worth (D/FW) Airport cooling system. However, characteristics of the systems studied by Lau (1985) and Hackner (1985) are also utilized. Table 1 summarizes the different component characteristics considered in this study. The loads employed throughout are representative of the D/FW system, so that components studied by Lau and Hackner are scaled for D/FW conditions. The ventilation air flow is taken to be 10% of the design air flow for the air handler under all circumstances. The use of economizer or "free" cooling cycles was not considered in this study. These modes of operation would occur at low ambient wet-bulb temperatures, such that the cooling loads could be met without operating the chillers.

The base system makes use of the component characteristics associated with the first choice for each component type in Table 1. For the most part, this corresponds to the current D/FW system. However, detailed data were not available for the performance of the air handlers. As a result, the air-handler performance characteristics from manufacturers' data used by Hackner (1985) were utilized in the base system. Except where otherwise noted, all pumps and fans (including air handlers) are considered to be operated with variable-speed motors. Although combinations of components from Table 1 other than that associated with the base system were considered, results are only presented for the base system, except when alternative systems yield different conclusions.

Many of the results presented in this paper are for steady-state conditions and are given as a function of load at ambient dry- and wet-bulb temperatures of 80°F and 70°F. For the purpose of performing simulations over a cooling season, optimal costs of operation are correlated in terms of the load and ambient conditions using the form given by Braun (1989). The result is integrated over time in response to time-varying load and weather conditions. In those cases where cooling season results are presented,

TABLE 1
Summary of System Component Characteristics

Component	Component Characteristics
Chillers	1) D/FW variable-speed chiller 2) D/FW fixed-speed chiller 3) Lau (1985) chiller scaled to D/FW loads 4) Hackner (1985) chiller scaled to D/FW loads
Cooling Towers	1) D/FW tower characteristics 2) Lau (1985) tower scaled to D/FW loads 3) Hackner (1985) tower scaled to D/FW loads
Pumping	1) D/FW system and pump characteristics 2) 20% greater system and pumping head at design conditions than 1) 3) 20% lower system and pumping head at design conditions than 1)
Air Handlers	1) Hackner (1985) scaled to D/FW loads 2) 20% greater fan power requirements at design conditions than 1) 3) 20% lower fan power requirements at design conditions than 1)
Loads	1) 20% of zone loads are latent 2) 15% of zone loads are latent 3) 25% of zone loads are latent

weather data for May through October in Dallas, Texas, and Miami, Florida, were utilized. Two different constant internal gains to the zones were considered: one-third and one-half of the maximum chiller cooling capacity. Only conditions where the ambient wet-bulb temperature was greater than 60°F were considered in the determination of plant operating costs.

CONTROL GUIDELINES FOR MULTIPLE COMPONENTS IN PARALLEL

For some of the independent control variables, it is possible to determine simple control guidelines that yield near-optimal performance. These guidelines simplify the optimization process in that these independent control variables may be reduced to dependent variables. They are also useful to plant engineers as "rules of thumb" for improved control practices. In this section, component modeling and optimization techniques are used to identify control guidelines associated with multiple components arranged in parallel.

Multiple Cooling Tower Cells

The power consumption of a chiller is sensitive to the condensing water temperature, which is, in turn, affected by both the condenser water and cooling tower air flow rates. Increasing either of these flows reduces the chiller power requirement, but at the expense of an increase in the pump or fan power consumption.

Braun (1987) and Nugent (1988) have shown that for variable-speed fans, the minimum power consumption results from operating all cooling tower cells under all conditions. The power consumption of the fans depends upon the cube of the fan speed. Thus, for the same total air flow, operating more cells in parallel allows for lower individual fan speeds and overall fan power consumption. An additional benefit associated with full-cell operation is lower water pressure drops across the spray nozzles, which results in lower pumping power requirements. However, at very low pressure drops, inadequate spray distribution may adversely affect the thermal performance of the cooling tower. Another economic consideration is the greater water loss associated with full-cell operation.

Most current cooling tower designs utilize multiple-speed, rather than continuously adjustable, variable-speed fans. In this case, it is not optimal to operate all tower cells under all conditions. The optimal number of cells operating and individual fan speeds will depend upon the system characteristics and ambient conditions. However, simple relationships exist for the best sequencing of cooling tower fans as capacity is added or removed. When additional tower capacity is required, then in almost all practical cases, the tower fan operating at the lowest speed (including fans that are off) should be increased first. Similarly, for removing tower capacity, the highest fan speeds are the first to be reduced.

These guidelines are derived from evaluating the incremental power changes associated with fan sequencing. For two-speed fans, the incremental power increase associated with adding a low-speed fan is less than that for increasing one to high speed if the following condition is satisfied.

$$\gamma_{t,low}^3 < (1 - \gamma_{t,low}^3) \tag{1}$$

or

$$\gamma_{t,low} < 0.79 \tag{2}$$

where $\gamma_{t,low}$ is the relative fan speed at low speed. If the low speed is less than 79% of the high fan speed, then the incremental power increase is less for adding a low-speed as opposed to a high-speed fan. In addition, if the low speed is greater than 50% of the high speed, then the incremental increase in air flow is greater (and therefore better thermal performance) for adding the low-speed fan. Most commonly, the low speed of a two-speed cooling tower fan is between one-half and three-quarters of full speed. In this case, tower cells should be brought on-line at low speed before any operating cells are set to high speed. Similarly, the fan speeds should be reduced to low speed before any cells are brought off-line.

For three-speed fans, the sequencing logic is not as obvious. However, for the special case where low speed is greater than or equal to one-third of full speed and the difference between the high and intermediate speeds is equal to the difference between the intermediate and low, then the best strategy is to increment the lowest fan speeds first when adding tower capacity and decrement the highest fan speeds when removing capacity. Typical three-speed combinations that satisfy this criterion are 1) one-third, two-thirds, and full speed or 2) one-half, three-quarters, and full speed.

Another issue related to control of multiple cooling tower cells having multiple-speed fans concerns the distribution of water flow to the individual cells. Typically, the water flow is divided equally among the operating cells. However, the overall thermal performance of the cooling tower is best when the flow is divided such that the ratio of water to air flow rates is identical for all cooling tower cells. In comparing equal flow rates to equal flow rate ratios, at worst, a 5% difference between the heat transfer effectivenesses for a combination of two tower cells, one operating at one-half and the other at full speed, was found. Depending upon the conditions, these differences generally result in less than a 1% change in the chiller power. These differences should also be contrasted with the lower water pressure drop across the spray nozzles (lower pumping power) associated with equally divided flow. In addition, the performance differences are smaller for greater than two-cell operation, when a majority of cells are operating at the same speed. Overall, equal water flow distribution between cooling tower cells is near-optimal.

Multiple Chiller Control

Multiple chillers are normally configured in a parallel manner and typically controlled to give identical chilled water supply temperatures. For the parallel chiller combinations considered in this study, controlling for identical set temperatures was found to be either optimal or near-optimal. Besides the chilled water setpoint, additional control variables are the relative chilled and condenser water flow rates. Simple guidelines may be established for distributing these flows.

In general, the relative condenser water flows to each chiller should be controlled to give identical leaving con-

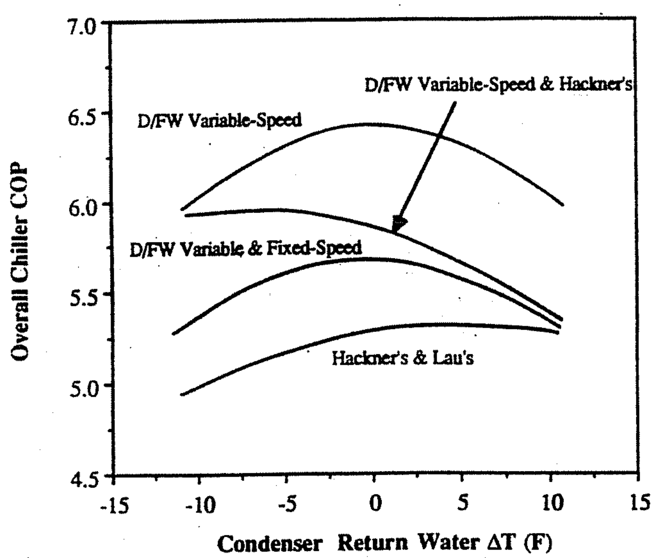

Figure 2 *Effect of condenser water flow distribution for two chillers in parallel*

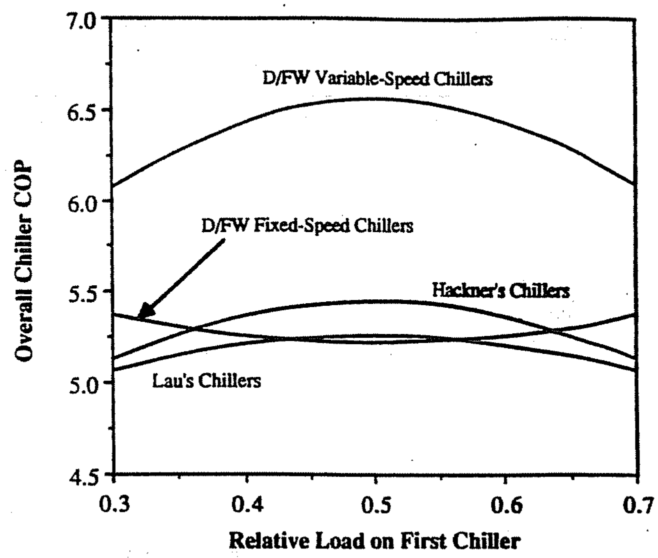

Figure 3 *Effect of relative loading for two identical parallel chillers*

denser water temperatures for all chillers. This condition approximately corresponds to relative condenser flow rates equal to the relative loads on the chillers. Figure 2 shows results for four different sets of two chillers operated in parallel. The overall chiller coefficient of performance (COP) is plotted vs. the difference between the condenser water return temperatures of the two chillers at equal loadings. For the identical D/FW variable-speed chillers, the optimal temperature difference is almost exactly zero. This was found to be the case for all identical chillers considered in this study. For the non-identical chillers of Figure 2, equal leaving condenser water temperatures result in chiller performance that is close to the optimum. Similar results were obtained for unequal loadings on the chillers.

For given chilled water return and supply temperatures, the relative chilled water flow to each chiller is equal to its relative loading. Consider the problem of determining the optimal relative loadings for N_{ch} chillers in parallel. The relative condenser water flows are assumed to be controlled to give identical return water temperatures. The optimization problem is one of minimizing

$$P_{ch} = \sum_{i=1}^{N_{ch}} P_{ch,i} \tag{3}$$

with respect to the relative loadings, $f_{L,i}$, with the constraint

$$\sum_{i=1}^{N_{ch}} f_{L,i} = 1 \tag{4}$$

By forming the Lagrangian and applying the first-order condition for a minimum or maximum, it can be shown that the point of minimum or maximum overall power occurs where the derivatives of the individual chiller power consumptions with respect to their relative loadings are equal.

$$\frac{\partial P_{ch,i}}{\partial f_{L,i}} = \frac{\partial P_{ch,j}}{\partial f_{L,j}} \qquad \text{for all i, j} \tag{5}$$

This condition, along with the constraint of Equation 4, is sufficient to determine the relative loadings. For identical chillers, these equations are satisfied for all chillers loaded equally or

$$f_{L,i} = \frac{1}{N_{ch}} \qquad \text{for i = 1 to } N_{ch} \tag{6}$$

For chillers with different cooling capacities, but identical part-load characteristics, the constrained optimality conditions are satisfied when each chiller is loaded according to the ratio of its capacity to the sum total capacity of all operating chillers. For the i[th] chiller,

$$f_{L,i} = \frac{Q_{cap,i}}{\sum_{i=1}^{N_{ch}} Q_{cap,i}} \tag{7}$$

where $Q_{cap,i}$ is the cooling capacity of the i[th] chiller.

The relative loadings determined with Equation 6 or 7 could result in either minimum or maximum power consumptions. With the second-order necessary condition for a minimizing point, it is possible to show that these points represent a minimum when the derivative of the COP with respect to relative loading is less than zero. In other words, the chillers are operating at loads greater than the point at which the maximum COP occurs. Typically, but not necessarily, the maximum COP occurs at loads that are about 40% to 60% of a chiller's cooling capacity. In this case, and with loads greater than about 50% of cooling capacity, the control defined by Equations 6 and 7 results in a minimum power consumption.

Figure 3 shows the effect of the relative loading on chiller COP for different sets of identical chillers loaded at approximately 70% of their overall capacities. Three of the chillers have maximum COPs when evenly loaded, while the fourth (D/FW fixed-speed) obtains a minimum at that point. The part-load characteristic of the D/FW fixed-speed chiller is unusual in that the maximum COP occurs at its capacity. This chiller was retrofit with a different refrigerant

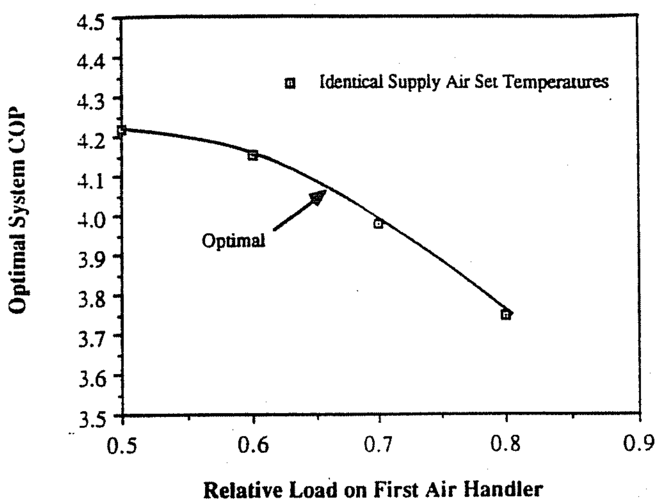

Figure 4 *Comparison of optimal system performance for individual supply air setpoints with that for identical values*

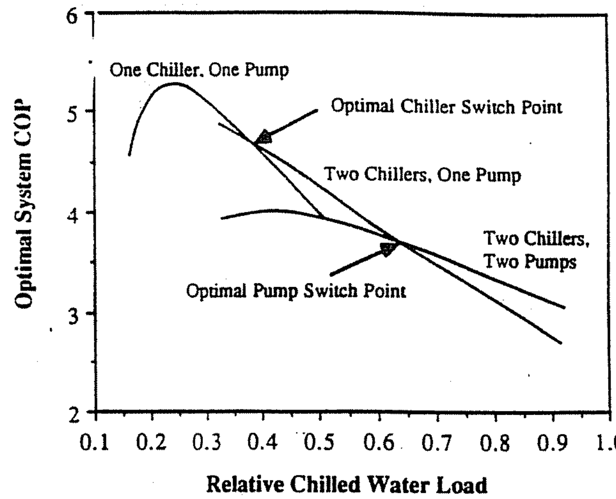

Figure 5 *Effect of chiller and pump sequencing on optimal system performance*

and drive motor, which caused its capacity to be derated from 8700 tons to 5500 tons. As a result, the chiller operates at a much smaller fraction of its original capacity.

The effect of loading on non-identical chillers was also investigated. The minimum power consumption was realized with near-even loading for all combinations considered, except for those involving the D/FW fixed-speed chiller. The best strategy for this particular fixed-speed chiller in combination with other chillers is to load it as heavily as possible, since its performance is best at full load.

One of the important issues concerning control of multiple chillers is chiller sequencing. Sequencing involves determining the conditions at which specific chillers are brought on-line or off-line. The optimal sequencing of chillers depends primarily upon their part-load characteristics. Chillers should be brought on-line at conditions where the total power of operating with the additional chiller would be less than without it. Optimization results indicate that the optimal sequencing of chillers may not be decoupled from the optimization of the rest of the system. The characteristics of the system change when a chiller is brought on-line or off-line due to changes in the system pressure drops and overall part-load performance. The optimal point for switching chiller operation may differ significantly from switch-points determined if only chiller performance were considered at the conditions before the switch takes place.

Multiple Air Handlers

In a large system, a central chilled water facility may provide cooling to several buildings, each of which may have a number of air-handling units in parallel. If the supply air setpoints of each of the air handlers were considered to be a unique control variable, the optimization problem would become quite complicated. However, the error associated with using identical supply air temperatures for all air handlers is relatively small as compared with the optimal solution, even when the loading on the various cooling coils differs significantly. As a result, the number of control variables in the optimization process may be reduced by one less than the number of air handlers.

Figure 4 shows a comparison between individual and identical setpoint control for a system with two identical air handlers. The system coefficient of performance associated with optimal control is plotted vs. the relative loading on one air handler. The difference between individual and identical setpoint control is not significant over the practical range of relative loadings. This result is most easily explained by considering the limits on the relative loadings. For the case of equal loadings on the air handlers, the optimal control setpoints are equal for identical air handlers. In the other extreme, where all of the load is applied to one air handler, the supply air temperature of the unloaded air handler has no importance, since its power consumption is zero. Between the two extremes, the error associated with assuming identical setpoints is relatively small. This result also extends to many air handlers in parallel and to non-identical designs.

Although the number of control variables is reduced by considering only a single supply air set temperature, the overall operating cost still depends upon the performance and loadings on the individual air handlers. However, for the purpose of determining optimal control, air handlers may be combined into a single effective air handler under the conditions that all zones are maintained at the same air temperature and the heat transfer characteristics of the coils are similar. The required air flow as a fraction of the design air flow for the i[th] air handler in parallel, $\gamma_{ahu,i}$, may be expressed in terms of an overall air flow ratio as

$$\gamma_{ahu,i} = \frac{f_{ahu,i}}{f_{ahu,des,i}} \gamma_{ahu} \tag{8}$$

where γ_{ahu} is the ratio of the total required air flow for all air handlers to the total design air flow, $f_{ahu,i}$ is the ratio of air flow for the i[th] air handler to the total air flow for all air handlers, and $f_{ahu,des,i}$ is the ratio of design air flow for the i[th] air handler to the total design air flow for all air handlers. If all air handlers have identical supply air temperature setpoints and zone air temperature setpoints, then $f_{ahu,i}$ is also

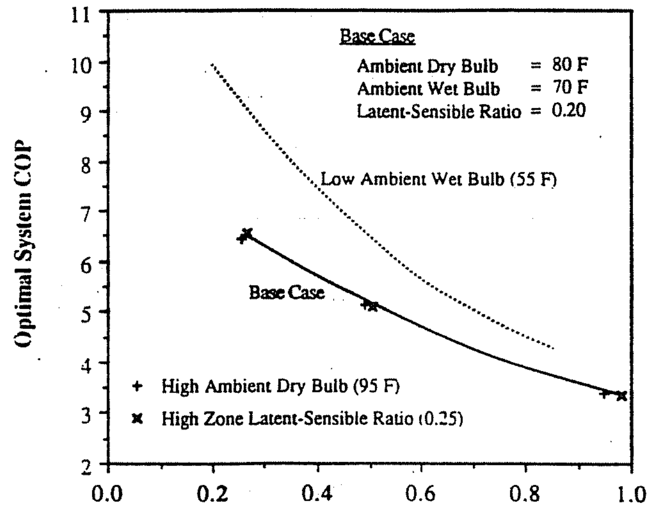

Figure 6 *Effect of uncontrolled variables on optimal system performance*

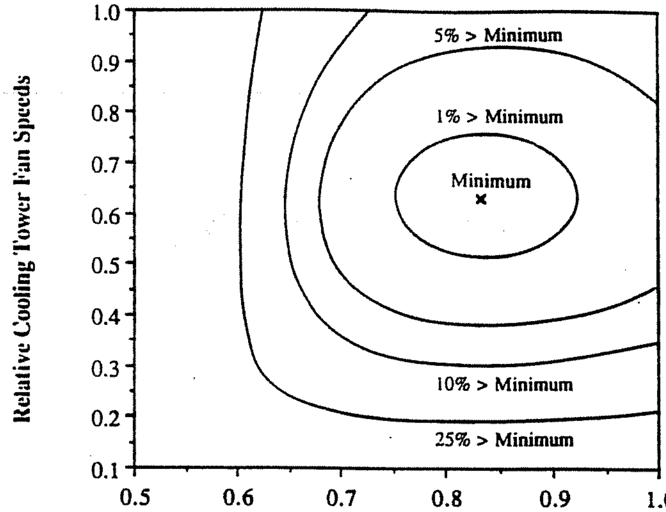

Figure 7 *Power contours for condenser loop control variables*

equal to ratio the sensible loading on the zones supplied by i^{th} air handler to the total sensible zone loads, $f_{sens,i}$. In this case, the total air handler power consumption for variable-speed fans is

$$P_{ahu} = \gamma_{ahu}^{3} \sum_{i=1}^{N_{ahu}} P_{ahu,des,i} \left[\frac{f_{sens,i}}{f_{ahu,des,i}} \right]^{3} \tag{9}$$

where $P_{ahu,des,i}$ is the power consumption for the i^{th} air handler at its design air flow.

As a result, all the air handlers may be combined into one air handler having the sum total area, water flow, air flow, and loading, with the power computed according to Equation 9. For the near-optimal control algorithm described in a companion paper (Braun 1989), it is necessary to include the relative zone sensible loads ($f_{sens,i}$) as uncontrolled variables in the empirical system cost function, if they change significantly.

Multiple Pumps

A common control strategy for sequencing both condenser and chilled water pumps is to bring pumps on-line or off-line with chillers. In this case, there is a condenser and chilled water pump associated with each chiller. For fixed-speed pumps, this strategy is not optimal. At the point at which a chiller is brought on-line in parallel and assuming that the pump control does not change, there is a reduction in the pressure drop and subsequent increase in flow rate for both the condenser and chilled water loops. The increased flow rates tend to improve the overall chiller performance. However, if the pumps were operating near their peak efficiency (a good design), then there is a drop in the pump efficiency when adding the additional chiller while holding the pump control constant. Most often, the improvements in chiller performance offset the degradation in pump performance, so that there is no need for an additional pump at the switch point.

Figure 5 shows the optimal system coefficient of performance (COP) for different combinations of chillers and fixed-speed pumps in parallel as a function of load relative to the total cooling capacity. The optimal switch point for a second pump occurs at a much higher relative load (~ 0.62) than the switch point for adding or removing a chiller (~ 0.38). If the second pump were sequenced with the second chiller, then the optimal switch would occur at the maximum capacity of one chiller and approximately a 10% penalty in performance would result at this condition. The optimal control for sequencing fixed-speed pumps depends upon the load and the ambient wet-bulb temperature and should not be directly coupled to the chiller sequencing.

The sequencing of variable-speed pumps is more straightforward than that for fixed-speed pumps. For a given set of operating chillers, variable-speed pumps have relatively constant efficiencies over their range of operation. As a result, the best sequencing strategy is to select pumps so as to operate near their peak efficiencies for each possible combination of operating chillers. Since the system pressure drop characteristics change when chillers are added or removed, the sequencing of variable-speed pumps should be directly coupled to the sequencing of chillers. For identical variable-speed pumps oriented in parallel, the best overall efficiency is obtained if they operate at identical speeds. For non-identical pumps, near-optimal efficiency is realized if they operate at equal fractions of their maximum speed.

SENSITIVITY ANALYSES AND CONTROL CHARACTERISTICS

In this section, the sensitivity of the optimal system performance to the uncontrolled and controlled variables is studied.

Effects of Load and Ambient Conditions

For a given system in which the relative loadings on each zone are relatively constant, the optimal control variables are primarily a function of the total sensible and latent gains for all zones, along with the ambient dry- and wet-bulb temperatures. Figure 6 shows the effect of these uncontrolled variables on the optimal system performance.

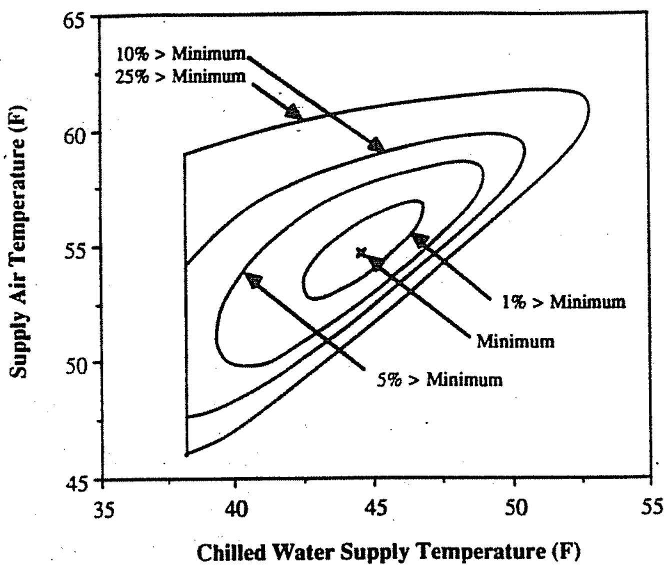

Figure 8 *Power contours for chilled water and supply air temperature*

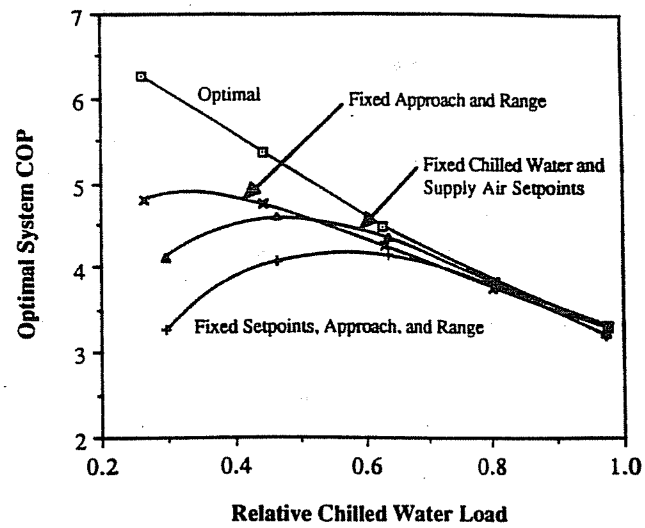

Figure 9 *Comparisons of optimal control with "conventional" control strategies*

For a given load and wet-bulb temperature, the effect of the ambient dry-bulb temperature is insignificant, since air enthalpies depend primarily upon wet-bulb temperatures and the performance of wet surface heat exchangers is driven primarily by enthalpy differences. Typically, the zone latent gains are on the order of 15%-25% of the total zone gains. In this range, Figure 6 shows that the effect of changes in latent gains has a relatively small effect upon the system performance for given total load. Consequently, results for overall system performance and optimal control may be correlated in terms of only the ambient wet-bulb temperature and total chilled water load. In the event that the load distribution between zones changes significantly through time, then this must also be included as a correlating variable.

Condenser Water Loop

The primary controllable variables associated with heat rejection to the environment are the condenser water and tower air flow rates. Both optimal air and water flows increase with load and wet-bulb temperature. Higher condensing temperatures and reduced chiller performance result from either increasing loads or wet-bulb temperatures for a given control. Increasing the air and water flow under these circumstances reduces the chiller power consumption at a faster rate than the increases in fan and pump power.

Figures 7 shows the sensitivity of the total power consumption to the tower fan and condenser pump speed. Contours of constant power consumption are plotted vs. fan and pump speeds. Near the optimum, power consumption is not sensitive to both of these control variables, but increases more quickly away from the optimum. The rate of increase in power consumption is particularly large at low condenser pump speeds. There is a minimum pump speed necessary to overcome the static head associated with the height of the water discharge in the cooling tower above the take-up from the sump. As the pump speed approaches this value, the condenser flow approaches zero and the chiller power increases dramatically. It is generally better to have too high rather than too low a pump speed. The "flatness" near the optimum indicates that it is not necessary for an accurate determination of the optimal control.

Chilled Water Loop

Both the optimal chilled water and supply air temperatures decrease with increasing load for a fixed wet-bulb temperature. This behavior occurs because the rate of increase in air handler fan power with respect to load increases is larger than the rate of decrease in chiller power at the optimal control points. As the wet-bulb temperature increases for a fixed load, the optimal set temperatures also increase. There are two primary reasons for this result: 1) For a given load, the chiller power depends primarily upon the temperature difference between the leaving condenser and chilled water temperatures. The condenser temperature and chiller power consumption increase with increasing wet-bulb temperature and the optimal chilled water temperature increases in order to reduce the temperature difference across the chiller. 2) In the absence of humidity control, optimal supply air and chilled water temperatures increase with decreasing sensible zone loads for constant total chiller load. In addition, the sensible to total load ratio decreases with increasing wet-bulb temperature for constant total load.

Figure 8 shows the sensitivity of the system power consumption to the chilled water and supply air set temperatures for a given load and wet-bulb temperature. Within about 2°F of the optimum, the power consumption is within 1% of the minimum. Outside this range, the sensitivity to the setpoints increases significantly. The penalty associated with operation away from the optimum is greater in the direction of smaller differences between the supply air and chilled water setpoints. As this temperature difference is reduced, the required flow of chilled water to the coil increases and the chilled water pumping power is larger. For a given chilled water or supply air temperature,

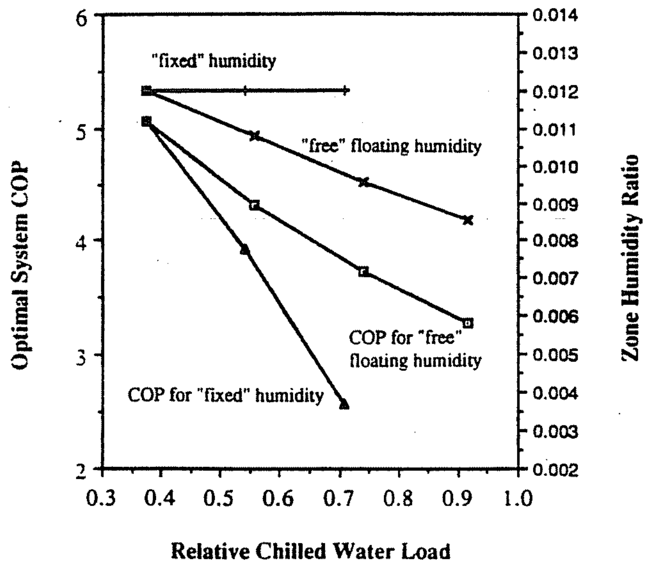

Figure 10 *Comparison of "free-floating" and "fixed" humidity control*

the temperature difference is limited by the heat transfer characteristics of the coil. Below this limit, the required water flow and pumping power would approach infinity if the pump output were not constrained. It is generally better to have too large rather than too small a temperature difference between the supply air and chilled water setpoints.

OPTIMAL VS. "ALTERNATIVE" CONTROL STRATEGIES

No general strategy has been established for controlling chilled water systems. Most commonly, the chilled water and supply air set temperatures are constants that do not vary with time. In some applications, these setpoints vary according to the ambient dry-bulb temperature. Generally, there is an attempt to control the cooling tower and condenser water flow in response to changes in the load and ambient wet-bulb. One strategy for controlling these flow rates is to maintain constant temperature differences between the cooling tower outlet and the ambient wet-bulb (approach) and between the cooling tower inlet and outlet (range), regardless of the load and wet-bulb. In some applications, humidity and temperature are controlled within the zones. In this section, the performance associated with some of these control strategies is compared with that of an optimally controlled system.

Conventional Control Strategies

Fixed values of chilled water and supply air setpoints and tower approach and range that are optimal or near-optimal over a wide range of conditions do not exist. In addition, it is not obvious how to choose values that work best overall. One simple, yet reasonable approach is to determine fixed values that result in near-optimal performance at design conditions. Figure 9 shows a comparison between the performance associated with optimal control and 1) fixed chilled water and supply air temperature setpoints (40° and 52°F) with optimal condenser loop control, 2) fixed tower approach and range (5° and 12°F) with optimal chilled water loop control, and 3) fixed setpoints, approach, and range. The results are given as a function of load for a given wet-bulb temperature. Since the fixed values were chosen to be appropriate at design conditions, the differences in performance as compared with optimal control are minimal at high loads. However, at part-load conditions, Figure 9 shows significant savings associated with the use of optimal control. Optimal control of the chilled water loop results in greater savings than that for the condenser loop at part-load ratios less than about 50%.

TABLE 2
Cooling Season Results for Optimal vs. Conventional Control

Control Description	Cost Relative to Optimal Control			
	Low Internal Gains		High Internal Gains	
	Dallas	Miami	Dallas	Miami
fixed temperature setpoints	1.09	1.10	1.03	1.03
fixed approach and range	1.05	1.06	1.03	1.03
fixed setpoints, approach, range	1.17	1.19	1.07	1.07

The overall savings over a cooling season for optimal control depends upon the variability of the load. If the cooling load was relatively constant, then fixed values of temperature setpoints, approach, and range could be chosen to give near-optimal performance. Table 2 summarizes cooling season operating costs for the conventional control strategies relative to optimal control for two different load characteristics in both Dallas and Miami. The low and high internal gains are approximately one-third and one-half of the maximum cooling requirement of the system. In Dallas, the maximum total operating cost difference of approximately 17% occurs for fixed setpoints, approach, and range with low internal gains. The difference is reduced to 7% for the high internal gains, since the system operates at a more uniform load near the design conditions. There is approximately twice the penalty associated with the use of fixed chilled water and supply air setpoints as compared with fixed tower approach and range for low internal gains, but the penalties are equal for the high internal gains. The results do not differ significantly for the Miami climate.

Humidity Control

In a variable-air-volume system, it is generally possible to adjust the chilled water temperature, supply air temperature, and air flow rate in order to maintain both temperature and humidity. In constraining the room humidity, the number of "free" control variables in the optimization is reduced by one. For a given chilled water temperature, there is at most one combination of the supply air temperature, air flow rate, and water flow rate that will maintain both the room temperature and humidity.

ASHRAE (1981) defines acceptable bounds on room temperature and humidity for human comfort. For a zone that is being cooled, the equipment operating costs are minimized when the zone temperature is at the upper bound of the comfort region. However, operation at the humidity upper limit does not minimize costs. Figure 10 compares costs and humidities associated with fixed (at the upper comfort limit) and free-floating zone

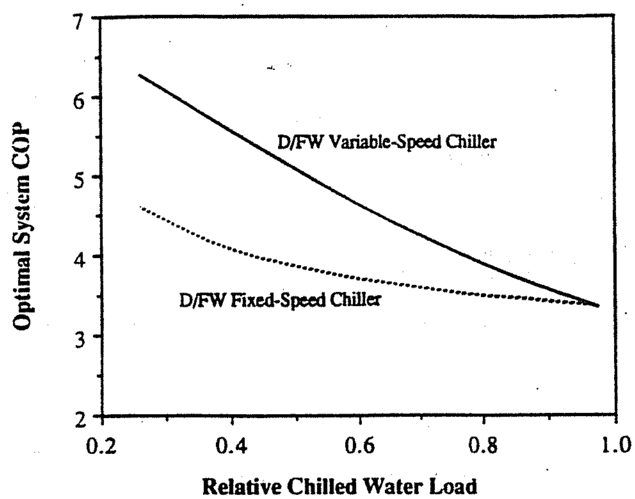

Figure 11 *Optimal system performance for variable- and fixed-speed chillers*

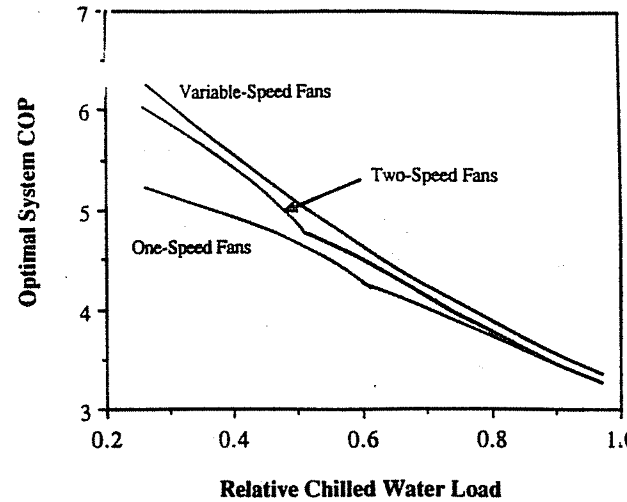

Figure 12 *Comparison of one-speed, two-speed, and variable-speed cooling tower fans (four cells)*

humidities as a function of the load. Over the range of loads for this system, the freefloating humidity operates with lower costs and humidities, with the largest differences occurring at the high loads. Operation at the upper humidity bound results in lower latent loads, but the addition of this humidity control constraint requires higher supply air temperatures than that associated with free-floating humidities. In turn, the higher supply air temperature results in greater air handler power consumption. In effect, the addition of any constraint reduces the number of free control variables by one and results in operation at a higher cost. In the determination of optimal control points, the humidity should be allowed to float freely, unless it falls outside the bounds of human comfort.

ALTERNATIVE SYSTEM CONFIGURATIONS UNDER OPTIMAL CONTROL

In order to compare the operating costs associated with different system designs, it is most appropriate that each be optimally controlled. In this section, alternative system configurations are compared in terms of their optimal system performance.

Variable vs. Fixed-Speed Equipment

The part-load performance of a centrifugal chiller depends upon the method by which its capacity is modulated. The Dallas/Fort Worth primary chiller was originally operated at a constant speed and the cooling capacity was modulated with the use of pre-rotation inlet vane and outlet diffuser vanes. The chiller was subsequently retrofit with a variable-speed electric motor and the vane control was disconnected.

In general, the part-load performance of a chiller is better for variable-speed as compared with fixed-speed control. The performance of the D/FW chiller was measured for both types of capacity modulation for nearly identical conditions, as described by Braun (1987). Figure 11 gives a comparison between the overall optimal system performance for both types of chiller control. At part-load conditions, the performance associated with the variable-speed control is significantly better. However, the power requirements are similar at conditions associated with the peak loads. This is expected, since, at this condition, the vanes are wide open and the speed under variable-speed control approaches that of the fixed-speed operation.

Part of the improvement with variable-speed chiller control may result from the unique characteristics of the D/FW chiller. The capacity of the D/FW chiller was derated so that the evaporator and condenser are oversized at the current capacity relative to the original design capacity. As a result, the performance is more sensitive to penalties associated with part-load operation of the compressor than to heat exchange improvements that occur with lower loads.

The most common design for cooling towers utilizes multiple tower cells in parallel that share a common sump. Each tower cell has a fan that may have one, two, or possibly three operating speeds. Although multiple cells having multiple fan settings offer great flexibility in control, the use of variable-speed tower fans can provide additional improvements in the overall system performance.

Figure 12 compares optimal system performance for single-speed, two-speed, and variable-speed tower fans as a function of load for a given wet-bulb. There are four tower cells for this system. All cells operate for the variable-speed fan results under all conditions, while cells are isolated for discrete fan control results when their fans are off. The discrete changes in the control of the multi-speed fans cause the discontinuities in the slopes of the curves in Figure 12. The flexibility in the control with one- or two-speed fans is most limited at low loads. Below about 70% of full-load conditions, the difference between one-speed and variable-speed fans becomes significant. With two-speed fans, the differences are on the order of 3% to 5% over the entire range.

If a fixed-speed pump is sized so as to give proper flow to a chiller at design conditions, then it is oversized for part-load conditions and the system will have higher operating costs than with a variable-speed pump having the same design capacity. The use of a smaller fixed-speed pump

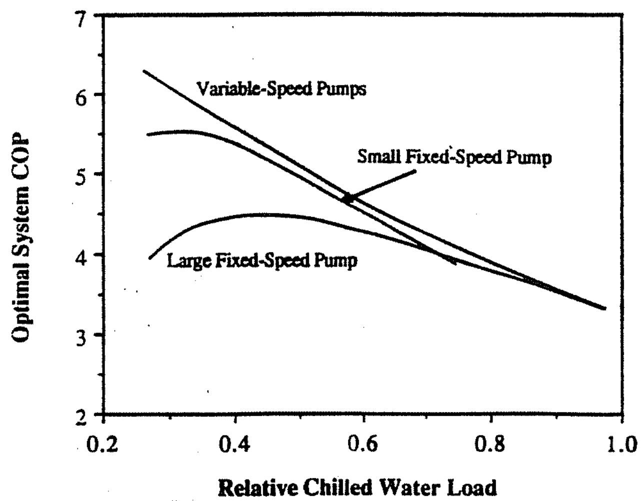

Figure 13 *Comparison of variable- and fixed-speed pumps*

TABLE 3
Cooling Season Results for Variable vs. Fixed-Speed Equipment

Configuration	Cost Relative to All Variable-Speed Equipment			
	Low Internal Gains		High Internal Gains	
	Dallas	Miami	Dallas	Miami
fixed-speed chiller	1.16	1.18	1.14	1.13
fixed-speed tower fans	1.06	1.06	1.05	1.05
fixed-speed pumps	1.10	1.09	1.05	1.04
fixed-speed air handler fans	1.07	1.07	1.04	1.04
all fixed-speed equipment	1.42	1.43	1.29	1.26

for low load conditions improves the flexibility in control and can reduce the overall power consumption. Figure 13 gives the optimal system performance for variable-speed and fixed-speed pumps applied to both the condenser and chilled water flow loops. The "large" fixed-speed pumps are sized for design conditions, while the "small" pump has one-half the flow capacity of the large. Below about 60% of full-load conditions, the use of variable-speed pumps shows a very significant improvement over single fixed-speed pumps. With the addition of "small" fixed-speed pumps, the improvements with variable speed become significant at about 40% of the maximum load.

The overall savings over a cooling season associated with the use of variable-speed equipment depends upon the variability in the load. Table 3 summarizes cooling season operating costs for the fixed-speed equipment relative to all variable-speed for low and high internal gains in both Dallas and Miami. Included in this table are results for fixed-speed air handler fans with variable-pitch blades to control the air flow. The results do not differ significantly between the Dallas and Miami climates. The largest differences for all situations are seen in comparing the fixed-speed to the variable-speed chillers (13% to 18%). The use of individual fixed-speed pumps or fans in the system carries a penalty of about 3% to 10% depending upon the load characteristics. Overall, the use of all-fixed-speed equipment results in operating costs that are 26% to 43% higher than for all variable-speed drives.

Series vs. Parallel Chillers

Multiple chillers are typically arranged in parallel and the chilled and condenser water flows are divided between the chillers according to their loading. Alternatively, it is possible to arrange chillers in series, so that the total chilled and condenser water flows pass through each evaporator and condenser. Two possible arrangements for series chillers are series-parallel and series-counterflow. In the series-parallel, the chilled and condenser water flows are in parallel, in that these streams enter the same chiller. For the series-counterflow arrangement, the streams enter at opposing chillers.

For the same total flow, multiple chillers operate more efficiently in series rather than in parallel. Figure 14 gives a comparison between the chiller coefficients of performance for parallel and series arrangements of two identical chillers as a function of the relative loading on the first chiller for identical entering temperatures and flow rates. Both series arrangements require significantly less power than the parallel chillers, the best arrangement being series counterflow. For the same total flow, the heat transfer coefficients are higher and leaving water temperature differences are lower for the individual chillers in series orientations than for parallel.

Although the chillers perform more efficiently in series rather than parallel, there are significant increases in water stream pressure drops across both the evaporators and condensers for the series arrangement. For two identical chillers, the ratio of the pressure drop across either the evaporator or condenser for a series arrangement as compared with that for parallel is approximately 8:1 for identical flow rates. The difference in the overall system performance for series vs. parallel depends upon the magnitude of evaporator and condenser pressure drops as compared with the other pressure losses in the chilled and condenser water loops. Figure 15 compares the optimal system performance for series and parallel chillers as a function of load for a given wet-bulb. For this system, the trade-offs between improved chiller performance and increased pressure drops balance such that overall system performance is similar for both configurations. Reducing the evaporator and condenser pressure drops by 20% did not have a significant effect upon this result.

SUMMARY

Optimization techniques were applied to analyzing the control of chilled water systems. The important uncontrolled variables that affect system performance and optimal control variable settings were identified as the total chilled water load and ambient wet-bulb temperature. Additional secondary uncontrolled variables that could be important if varied over a wide range would be the individual zone latent to sensible load ratios and the ratios of individual sensible zone loads to the total sensible loads for all zones.

Control simplifications that reduce the number of independent control variables and simplify the optimization were also identified. Using these general results, Braun (1989), in a companion paper, has presented a "simple" methodology for near-optimal control of chilled water systems without storage. The simplifications are also useful

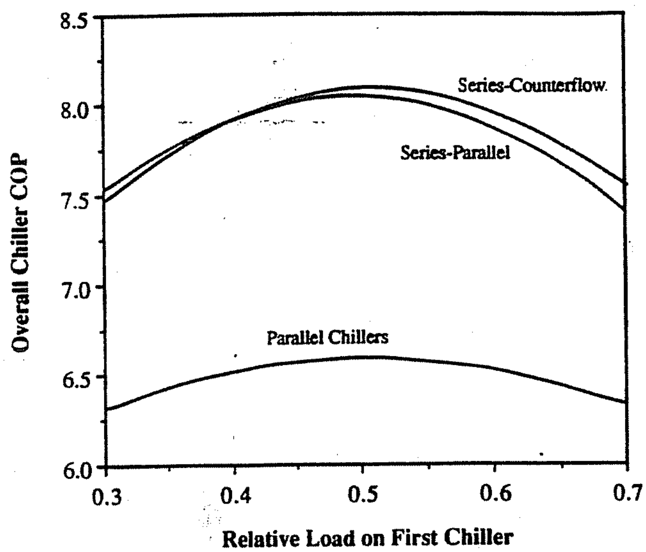

Figure 14 *Chiller performance for parallel and series chillers*

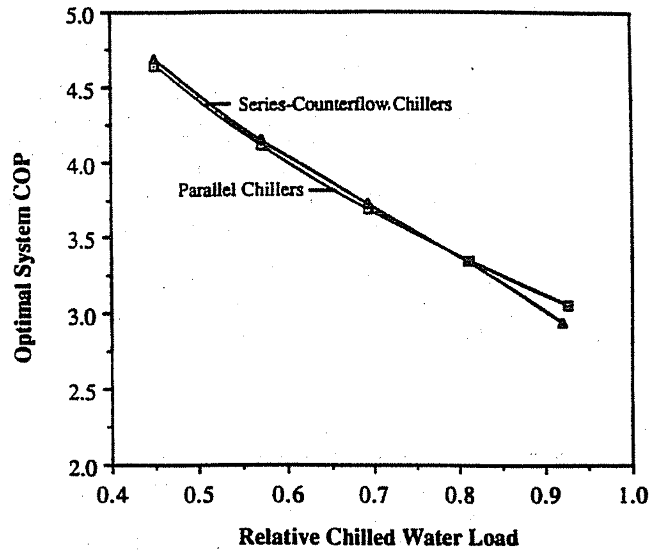

Figure 15 *Optimal system performance for series and parallel chillers*

to plant engineers for improved control practices and are summarized as follows:

1. *Variable-Speed Tower Fans:* Operate *all* tower cells at identical fan speeds.
2. *Multi-Speed Tower Fans:* Increment lowest-speed tower fans first when adding tower capacity. Reverse for removing capacity.
3. *Variable-Speed Pumps:* The sequencing of variable-speed pumps should be directly coupled to the sequencing of chillers to give peak pump efficiencies for each possible combination of operating chillers. Multiple variable-speed pumps should be controlled to operate at equal fractions of their maximum speed.
4. *Chillers:* Multiple chillers should have identical chilled-water set temperatures and the evaporator and condenser water flows for multiple chillers should be divided according to the chillers' relative cooling capacities.
5. *Air Handlers:* All parallel air handlers should have identical supply air setpoint temperatures.

No general simplifications could be found for the optimal sequencing of chillers and fixed-speed pumps. It is necessary to evaluate the overall system performance in order to determine the optimal points for adding or removing chillers. In general, it is not optimal to sequence fixed-speed pumps with chillers.

Additional results and conclusions concerning both control and design under optimal control of chilled water systems are summarized as follows:

1. Depending upon the load characteristics, fixed values of chilled water and supply air setpoints and cooling tower approach and range resulted in approximately 7% to 19% greater cooling season operating costs than that for optimal control in Dallas and Miami.
2. In the determination of optimal control points, the humidity should be allowed to float freely, unless it falls outside the bounds of human comfort. In effect, the addition of any constraint reduces the number of free control variables by one and results in operation at a higher cost.
3. Depending upon the load characteristics, the cooling season operating costs were approximately 26% to 43% greater for all fixed-speed equipment as compared with all variable-speed equipment in Dallas and Miami. The most significant difference was attributed to the chiller.
4. The performance of multiple chillers is enhanced by orientation in series rather than parallel. However, the increase in pumping power requirements for series chillers offsets the chiller improvements and the overall performance for the two configurations is similar.

REFERENCES

Arnold, J.A., et al. 1984. "Supervisory control and system optimization of chiller-cooling tower combinations via simulation using primary equipment models—part II simulation and optimization." Proceedings of the Workshop on HVAC Controls and Modeling and Simulation, Georgia Institute of Technology, February 2-3.

Braun, J.E. 1988. "Methodologies for the design and control of central cooling plants." Ph.D. Thesis, University of Wisconsin-Madison.

Braun, J.E.; Klein, S.A.; Beckman, W.A.; and Mitchell, J.W. 1989. "Methodologies for optimal control of chilled water systems without storage." ASHRAE Winter Annual Meeting, Chicago.

Braun, J.E.; Mitchell, J.W.; Klein, S.A.; and Beckman, W.A. 1987. "Performance and control characteristics of a large central cooling system." *ASHRAE Transactions*, Vol. 93, Part 1.

Hackner, R.J.; Mitchell, J.W.; and Beckman, W.A. 1984. "HVAC system and energy use in existing buildings—part I." *ASHRAE Transactions*, Vol. 90, Paper KC-84-09.

Hackner, R.J.; Mitchell, J.W.; and Beckman, W.A. 1985a. "HVAC system dynamics and energy use in buildings—part I." *ASHRAE Transactions*, Vol. 91, Part 1.

Hackner, R.J.; Mitchell, J.W.; and Beckman, W.A. 1985b. "System dynamics and energy use." *ASHRAE Journal*, June.

Johnson, G.A. 1985. "Optimization techniques for a centrifugal chiller plant using a programmable controller." *ASHRAE Transactions*, Vol. 91, Part 2.
Klein, S.A.; Nugent, D.R.; and Mitchell, J.W. 1988. "Investigation of control alternatives for a steam turbine driven chiller." *ASHRAE Transactions*, Vol. 94, Part 1. pp. 627-643.
Lau, A.S.; Beckman, W.A.; and Mitchell, J.W. 1985. "Development of computer control routines for a large chilled water plant." *ASHRAE Transactions*, Vol. 91, Part 1.
Marcev, C.L. 1980. "Steady-state modeling and simulation as applied to energy optimization of a large office building integrated HVAC system." M.S. Thesis, University of Tennessee.
Sud, I. 1984. "Control strategies for minimum energy usage." *ASHRAE Transactions*. Vol. 90. Part 2.

DISCUSSION

Z. Cumali, Principal, CCB/Cumali Associates, San Francisco, CA: Some of the conclusions presented in this paper do not agree with the results we have observed in our work.

The first impression we wish to correct is the fact that solution of optimal system operation problems does not require drastic specification of models to be applied in real time.

The building in which the simplified rules will result in more energy consumption is described in the reference paper. In summary the system has four cooling towers, three chillers and corresponding pumps, and four major variable volume systems. All units have variable speed control except the chillers.

1. Setting all supply air temperatures the same does not lead to even near-optimal results when the coil characteristics and loads are quite different from unit to unit; when any of the flows reach minimum or maximum on the air or water side; or if the relative cost of pumping water is larger than moving the air. An example of minimum air condition in one unit shows that

					Power Input
Load in Tons	117	42	75	83	kW
Supply air temperature, °F					
Optimal	52	59	52	52	216
Based on rule	59	59	59	59	247

If all temperatures are reduced to 52°F, then the second unit will have less than minimum required air or result in significant overcooling. The 12.6% change is typically what one expects to save from optimization.

2. Setting chilled water temperatures to be the same does not produce near-optimal results again when the chiller characteristics are different and minimum or maximum flow conditions are reached. Using the same cases as above but running two chillers, one large and old and one small and new, we have the following results:

Chilled water temperatures		Power input kW
42	54	343
42	42	363
49	54	375

In the last case we have reached the maximum flow condition in the first chiller and therefore have to limit the supply temperature to 49°F.

As these examples show, the conclusions stated in the two papers quite often result in considerable increases in energy input. It is therefore very important that the authors emphasize the conditions and the assumptions which significantly limit the applicability of their results. e.g., the similarity of performance characteristics of equipment and effects of operational constraints, etc.

Unfortunately, this level of simplification appears to be counterproductive and possibly misleading in that the casual reader is left with the impression that the problem may be solved with a few simple rules. Quite to the contrary, this field of study is complex and will require much research and many more Ph.D. theses from people of the caliber of Dr. Braun and his advisors, who are to be commended and encouraged for continuing the work presented in these papers.

Reference: "Global Optimization of HVAC System Operations in Real Time," Zulfikar Cumali, ASHRAE Transactions, presented in Dallas, TX, 1988.

J.E. Braun: We agree that the "best" solution for determining the optimal control for a given system is to have a detailed model of the complete process that operates in parallel with the actual system. An optimization algorithm is then applied to this model in order to determine the optimal control. The component-based optimization methodology presented in our paper addresses this goal. However, this type of aapproach requires detailed measurements for each component within the system in order to update parameters of the models so that they adequately match the real performance. Often these measurements are not available or are inaccurate. In addition, the description of the system to be modeled and optimized requires considerable expertise and the effort is beyond the capabilities of most installers of Energy Management Control Systems.

Our results indicate that this level of effort is not necessary in order to accomplish near-optimal control of these systems. Our paper develops a set of heuristic rules for good control, along with a system-based approach for determining near-optimal control set points. Mr. Cumali brings up an important point concerning the effects of operational constraints on the applicability of broad-based "rules of thumb." However, common sense heuristics for handling these constraints should give near-optimal results under most circumstances. Mr. Cumali questions two of the rules that were established and provides examples that he believes contradict our results. However, he appears to have misinterpreted and misapplied these rules.

The first example presented by Mr. Cumali concerns the rule of utilizing identical supply air set temperatures for all air handlers with variable-air-volume (VAV) control. The results of his optimal control analysis show three out of four air handlers operating with identical set temperatures with the fourth set at the minimum value necessary to maintain the minimum flow requirement. He then compares the power requirements for this case with setting all set temperatures equal to the one for the air handler that is constrained at its minimum flow. This is not a correct comparison. The rule does not state that the supply air temperatures should be set to the value for minimum flow air handler. This rule provides a simplification for determining a single optimal discharge air temperature for air handlers with VAV control, rather than having to treat all discharge air temperatures as separate variables. When an air handler is at minimum flow, then it is no longer utilizing VAV control. The flow rate is constant and the discharge air temperature should be adjusted to maintain the room condition. The discharge air temperatures of the remaining VAV controlled air handlers may be considered identical in the optimization. It is interesting to note that the optimal values of discharge air temperature for air handlers under VAV control in Mr. Cumali's example are identical.

The system-based methodology should be applied assuming that no constraints are violated. However, only those air handlers that are not operating at minimum flow should have their set points adjusted to the optimized value. Correct application of this procedure will give results that are much closer to the optimum than that for Mr. Cumali's example.

In Mr. Cumali's second example, he compares the optimal power consumption with that for identical chilled water set temperatures when operating two chillers. However, the basis for these comparisons is incorrect. Mr. Cumali compares his optimal (42 F and 54 F) with 1) identical set points equal to his lower optimized value of 42 F and 2) identical set points equal to his upper optimized value of 54 F (although the first chiller was constrained to operate at 49 F). The choice of identical set points for the comparisons is incorrect in that they are not optimized values. Each of the three cases result in significantly different overall chilled water supply temperatures. The advantage of utilizing identical set temperatures is that the optimization process is simplified. The value of this single set point cannot be arbitrarily established but must be estimated utilizing an optimization methodology (e.g., as outlined in our paper). For Mr. Cumali's example, an optimal single setpoint would fall somewhere between 42 F and 54 F. In order to handle flow constraints, the system-based methodology should be applied assuming that no constraints are violated. However, only those chillers that are not operating at minimum (or maximum) flow should have their set points adjusted to the optimized value.

Mr. Cumali has pointed out that a detailed approach yields reductions in power consumption over simplified procedures. However, this requires a considerable investment of time and money. We disagree with Mr. Cumali's assertion that the level of simplification of our methodology is counterproductive. This field has not evolved to a point where on-line optimal control of chilled water systems using detailed system models is widely applied. Appropriate application of the heuristics and simplified methodology developed for near-optimal control in this study provide significant improvements over current practice.