

Parameter Estimation and Optimal Control of Chilled Water Systems

Barrett Flake, John Mitchell and William Beckman

Abstract

The energy costs of operating chilled water cooling systems can be minimized through optimal control. A general approach for determining the optimal control involves the following tasks: 1) developing phenomenological models for the system components, 2) determining model parameters using measured data, and 3) subjecting the system of parametric models to an optimization algorithm. Parameter estimation is one of the most difficult tasks because of discontinuous variables and nonlinear relations between input and output variables. In addition, critical inputs for components are not always measured.

A physical plant was studied that consisted of interconnected components including an electric driven chiller, a chiller driven by a steam turbine, and a multi-cell cooling tower. Various methods for parameter estimation and control optimization were studied. These are evaluated using simulated and physical systems. Optimal supervisory control is determined through application of the simulated annealing method to the model. Cost savings of optimal control over conventional control strategies are compared.

1. Introduction

Large chilled water plants employ multiple chillers, cooling towers, pumps and other equipment to meet required chilled water loads. With the multiple components, each having discrete or continuously variable levels of operation, a plant can usually meet any given load over a wide range of possible operating conditions. For example, the chilled water load could possibly be met by operating a single chiller at near full load, or by operating two chillers, each at partial load. In addition to the many possible combinations of chiller loading, the chillers can usually be operated within a range of condenser water temperature. Thus, a particular load can be met over a range of cooling tower operation. Of all possible combinations of component operating levels, there exist one or more combinations that minimize energy costs.

2. Plant Description

A chilled water plant on the University of Wisconsin-Madison campus served as a model for applying methods of parametric estimation and optimal supervisory control. A schematic representation of one section of the plant is shown in Figure 1. This section of the plant includes a steam driven chiller (#3) and an electric motor driven chiller (#4). The design load is 5,500 tons for both chillers.

Chillers #3 and #4 are served by a two cell cooling tower with two two-speed fans (off/low/high) per cell. The two cells have a common sump. Chilled water pumps (CHP) draw from the chilled water return (chwr) header, pumping the water through the evaporative heat exchanger and into the chilled water supply (chws) header. The chilled water pumps are designed primarily to meet the pressure drop through the evaporator rather than provide a pressure differential between the main distribution supply and return chilled water line. Large main distribution pumps (not shown on the schematic) maintain pressure between the supply and return lines sufficient for distribution throughout the campus.

Condenser water pumps (CWP) circulate water between the refrigerant condenser and the cooling towers. For the steam driven turbine, a "surface condenser" is placed between the refrigerant condenser and the cooling tower. This heat exchanger is used to condense the steam turbine exhaust before returning the condensate to the boilers. Turbine exhaust pressures are normally subatmospheric with a saturated steam temperature around 37°C.

3. Parameter Estimation Approach

A transient system simulation program, TRNSYS, developed at the University of Wisconsin Solar Energy Laboratory, was used to model the chilled water system. The program is modular so that a number of individual component models may be linked to form a system. Many TRNSYS component models incorporate mechanistic design equations derived from first principles of heat or mass transfer and thermodynamics. Individual TRNSYS component models of chillers, heat exchangers and a cooling tower are linked in a TRNSYS "deck" of the Walnut Street Plant, where the connections between outputs of one component and inputs of others are defined.

The TRNSYS chiller component model predicts compressor power consumption, P_{comp} , and leaving condenser water temperature, T_{cwr} . Input variables to the chiller model include entering chilled water temperature, T_{chwr} , a set point for the leaving chilled water temperature, T_{chws} , entering

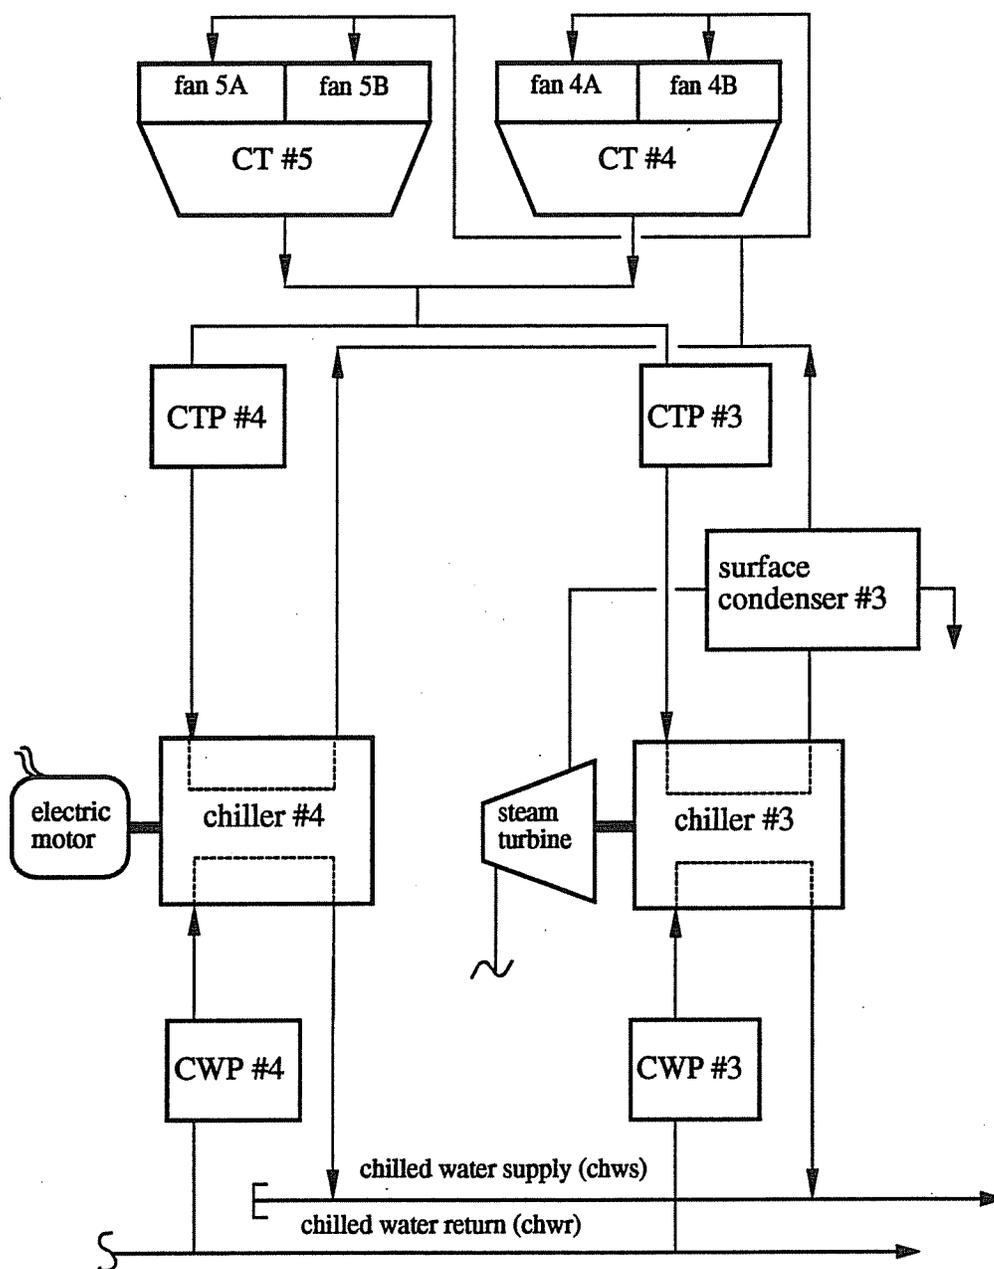


Figure 1: A schematic of the Walnut Street chilled water plant showing chilled wacondenser water flows.

condenser water temperature, T_{cws} , and both chilled and condenser water flow rates \dot{m}_{chw} , and \dot{m}_{cw} . A quadratic model is used in which six parameters relate compressor power to load and temperature difference between leaving condenser and chilled water flows. The sum of compressor energy and load is used along with condenser flow rate and specific heat to determine leaving condenser water temperature.

The steam turbine component model uses input and exhaust pressures, along with required power from the chiller component, to determine steam flow rate and exhaust enthalpy. Turbine power is calculated using an isentropic efficiency modeled as a linear function of pressure ratio and turbine power.

An effectiveness-NTU relationship is used to model the heat transfer in the condenser. The overall heat transfer coefficient is assumed to be constant since the condenser water flow is nearly constant and the condensing temperature does not vary widely. The TRNSYS cooling tower component model incorporates an effectiveness/Ntu relationship to model heat and mass transfer transport. [Bra88] A more detailed description of the component models including equations are given a previous report. [Fla97]

Dynamic response in the local control loops can be neglected in the optimization problem under certain assumptions. Hackner [Hac85] demonstrated that for systems without significant thermal storage, the system dynamics of local loop controls can be neglected in the determination of optimal supervisory controls. The chiller plant model used in this work is steady state with time varying external or forcing variables (e.g. chilled water return temperature and ambient wet bulb temperature.)

Ordinary least squares and the determinant criterion were used to formulate the regression problem for parameter estimation. Residuals corresponding to various combinations of responses (output dependent variables) were minimized in comparative parameter estimation runs. Ten measured responses are available from the physical plant: \dot{m}_{steam} , p_{exh} , T_{cdwi} , T_{cdwo} , T_{hw} , $T_{\text{cwi} \#3}$, $T_{\text{cwo} \#3}$, $T_{\text{cwi} \#4}$, $T_{\text{cwo} \#4}$, $P_{\text{chiller} \#4}$. Parameter estimates are needed for a total of 22 parameters. A detailed report of the parameter estimation problem and methods used for solution are found in an earlier report. [Fla97]

Data were available over a one month period, recorded every two hours. A best set of parameter estimates were found that gave a steam flow residual standard deviation of 867 kg/hr [8.3%] and electric power residual standard deviation of 127 kW [5.9%]. The better fit of electric chiller power was possible due to the availability of measured data for all input and output variables of the electric chiller model.

Given parametric models representative of HVAC system components and a means for solving the system of equations, the optimization problem can then be addressed. A general minimization problem statement for a system of n energy consuming components is given in Equation 1 below. The

quantity to be minimized, J , is the total cost of energy, summed over all components, for a given time period. The power, p_i , for each of the i components depends upon the controlled variables, c_i , and uncontrolled variables u_i . Energy sources for each component have associated unit costs, $k_i(t)$, which can be time dependent (e.g. time of day electrical rates). A set of controlled variables, c_i , is sought which minimizes the objective function, J .

$$\underset{c_i}{\text{minimize}} \quad J = \sum_{i=1}^n k_i p_i(c_i, u_i) \Delta t \quad (1)$$

Although the objective function is a linear combination of certain outputs (e.g. component power), the constraints will, in general, be nonlinear, and can be discontinuous. Also, some variables may only take on discrete or integer values. For example, some fans or pumps may have multiple speeds (e.g. off/low/med/high). For the general case, it is possible for the objective function to have multiple local minima, which restricts the number of available solution methods. A global optimization method called simulated annealing has demonstrated success in solving both combinatorial problems (e.g. mixed integer problems) and functions of continuous variables and was chosen for this problem.

A simulated annealing subroutine written by Goffe *et al.* [Gof94] was modified and incorporated into the TRNSYS simulation package. For given input variables and parameters, TRNSYS attempts to converge upon a stable set of outputs. The component models are not generally represented by continuous functions, but are FORTRAN subroutines which may include a number of logical statements that could make the output discontinuous or nonsmooth. Also, there may be constraints built into the subroutines such that outputs never exceed some upper or lower bounds. The discontinuities and nonsmooth relationships between component inputs and outputs limit the choice of optimization algorithms to only the most robust, such as simulated annealing.

4. Optimal Supervisory Control Problem

The objective is to minimize the system energy costs for the two chillers (one steam driven and one electrically driven), four cooling tower fans, two chilled water pumps and two condenser water pumps. Steam cost is estimated by physical plant personnel to be \$7.74/1000 kg. Electrical costs are dependent upon the time of day of usage. During "peak rate" usage hours (10:00 a.m. through 9:00 p.m. weekdays) the cost is \$0.0440/kW·hr, and the cost is \$0.0264/kW·hr during the "off-peak"

hours. An additional electrical demand charge of \$7.00/kW is levied for the highest power demand during on-peak hours in the one month billing period.

Certain loads served by the chilled water plant have a high limit on chilled water supply temperature. The water temperature from the plant is constrained to not exceed the limiting temperature of 4.6 °C. For example, if water flow rates through the individual chillers were equal, a chilled water temperature of 5.6 °C from one chiller would constrain the temperature from the other chiller to no greater than 3.6 °C. The supply temperature from one chiller is thus related to the supply from the other chiller, and this effectively removes one chilled water set temperature as a control variable.

The four cooling tower fan motors can individually be run at zero, half, or full speed. Since fan power increases in cubic proportion to fan speed, the change in power from half to full speed is greater than the power required at half speed. Also, the increase in cooling tower effectiveness is less from half to full speed than from off to full speed. Thus, the combinations having a fan off simultaneously with a fan on at full speed will always use more power and result in less cooling than both fans at half speed. Since all four cells of the cooling tower have equivalent performance, the particular combination or order of operating cells makes no difference. This allowed a single control variable having eight discrete values between 0 and 1 to represent the particular combinations of fan speed.

The resulting optimization problem is a nonlinear function of one continuous variable ($T_{\text{chws } \#3}$) and one discrete variable (fan speed control). The continuous control variable, chilled water supply temperature for chiller #3, $T_{\text{chws } \#3}$, has a lower bound of 3.3 °C (lowest safe operating temperature according to plant operators). The upper bound for $T_{\text{chws } \#3}$ is the warmest water temperature that forces the chilled water temperature out of chiller #4 to its minimum (also 3.3 °C) in order to meet the required combined flow temperature.

Conventional Control

A sequence of 31 days of data, beginning with July 1, is established as a period for comparing optimal control strategies with the actual control. Costs for conventional operation were estimated using the Walnut Street Plan model and the measured data for the period. In the simulation, the chilled water supply temperatures from both electric and steam driven chillers were set to match the average measured values over the period.

The operators' existing strategy for operating the plant is to supply chilled water between 4.4 °C and 6.1 °C. Cooling tower fans are manually controlled by the operators. Generally, fans are added or fan speed increased if the water temperature from the cooling tower exceeds about 21.1 °C. A simulation of the plant operating at the measured values of control variables (conventional control) is used as a basis for comparing the costs under optimal control. The steam, electrical and total energy costs over the 31 days are \$63,444, \$95,034 and \$158,478 respectively.

Optimal Supervisory Control

In the optimization calculations, values of the two control variables ($T_{\text{chws \#3}}$ and γ_{ct}) that minimize energy (steam and electricity) costs for each hour are found. The chilled water temperature from the electric chiller ($T_{\text{chws \#4}}$) is constrained such that the mixed water temperature is at the set point of 4.6 °C (40.3 °F). Over most of the period, optimal control favored using the electric chiller (#4) for the majority of the load.

In Figures 2 and 3 optimal chilled water supply temperatures from chiller #3 are plotted against chilled water return temperature for two values of electricity costs. The minimum total cost occurs with either the steam chiller or electric chiller fully loaded. At lower return water temperatures (less than 10.4 °C), the electric chiller is run at its minimum allowable temperature (fully loaded) and the steam chiller takes on the remainder of the load. At the higher electric rate it is optimal to shift from fully loading the electric chiller to fully loading the steam chiller at a load corresponding to a return temperature of 10.4 °C. For the lower electric rate it is optimal to fully load the electric chiller up to the maximum capacity. The reason for the control lies in the part load performance of the chillers. The performance (COP) of the steam driven chiller increases significantly with load so that it is advantageous to shift from the electric to the steam system at some load when electric rates are high. However, for low electric rates, the electric chiller is always cheaper to operate.

A plot indicating optimal cooling tower fan settings at various combinations of chilled water return and wet bulb temperatures at the lower electrical rate is given in Figure 4. The overlapping groupings of particular fan speed settings demonstrate dependence upon both return and wet bulb temperature. When operating at the higher electrical rate, switching to a higher fan speed setting occurs at a slightly warmer chilled water return temperature.

The operators' cooling tower fan control decisions are also shown as a function of wet bulb and chilled water return temperature in Figure 5. Although scattered, the operators' settings somewhat coincide with the optimal values. When asked how he decided when to increase fan speed, a senior operator at the Walnut Street Plant revealed that he could tell from the sound and vibration of a chiller when a cooler condenser temperature is needed.

With optimal control, the steam and electricity costs are reduced to \$52,770 and \$99,930 respectively. The total energy cost for the month is \$152,700 for a savings of \$5,778 (3.6%) over conventional control. Average savings per day is \$186.49.

The previous results were obtained without considering electrical demand charges. The power used in operating the electric chiller at full load during a peak period represents approximately \$25,000 in monthly demand charges.

The peak electrical demand under conventional control is 5,439 kW, corresponding to \$38,074 in peak demand charges. The peak demand in the optimal control without regard to demand is 5,646 kW, an increase of \$1,449 in demand costs. The savings in optimal over conventional control including demand costs is reduced to \$3850 for the month, 2.0% of the total energy and demand costs under conventional control. Almost half of the optimal savings are offset by increased demand charges. To find the minimum monthly energy cost for the chilled water plant, the demand charge must be included in the optimization.

One strategy for reducing peak demand costs is to minimize electricity use during an anticipated high cooling load period. For example, at the Walnut Street Plant, electrical energy could be reduced by moving as much of the load as possible to the steam driven chillers, unloading the electrically driven chiller. Knowing when a peak load occurs impacts the operators' decisions in loading the chillers. If a large peak was to occur later in the billing period, or if the peak had already passed, operators would not take the peak charges into account when operating the plant during periods of lower loads. But the true peak load for the billing period is unknown until the end of the billing period, and operators must make predictions of future load based on expected weather conditions.

In an actual implementation, determining an optimal control strategy to minimize both energy and demand costs requires weather and load predictions for the billing period. Using historic mea-

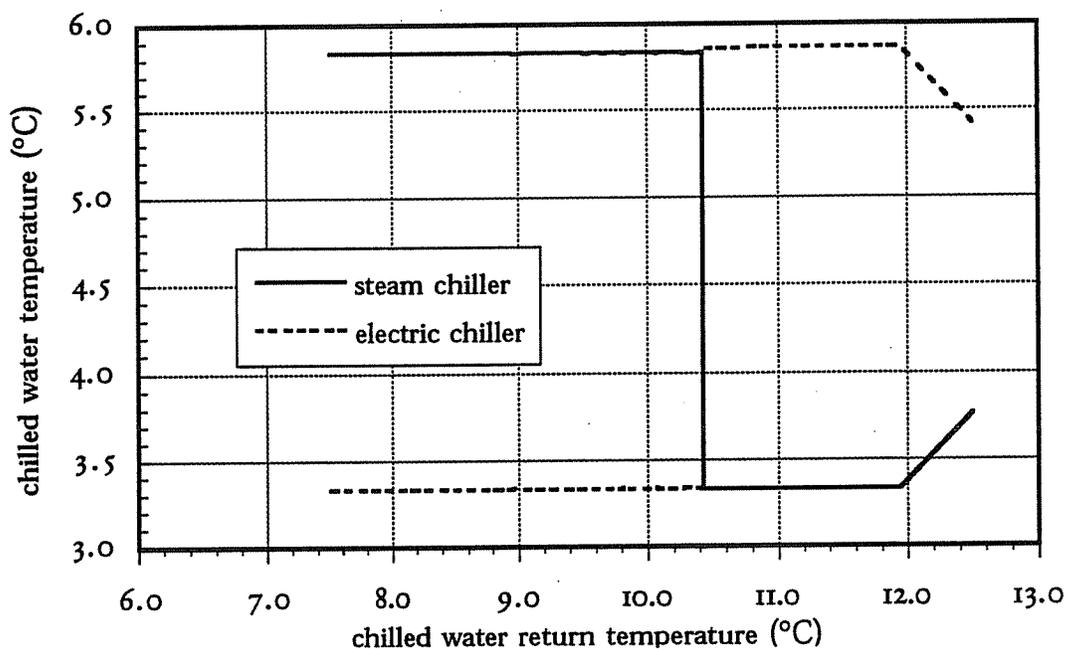


Figure 2: Optimal chilled water temperatures from both chillers as a function of chilled water return temperature. Electric costs are 4.40¢/kW·hr.

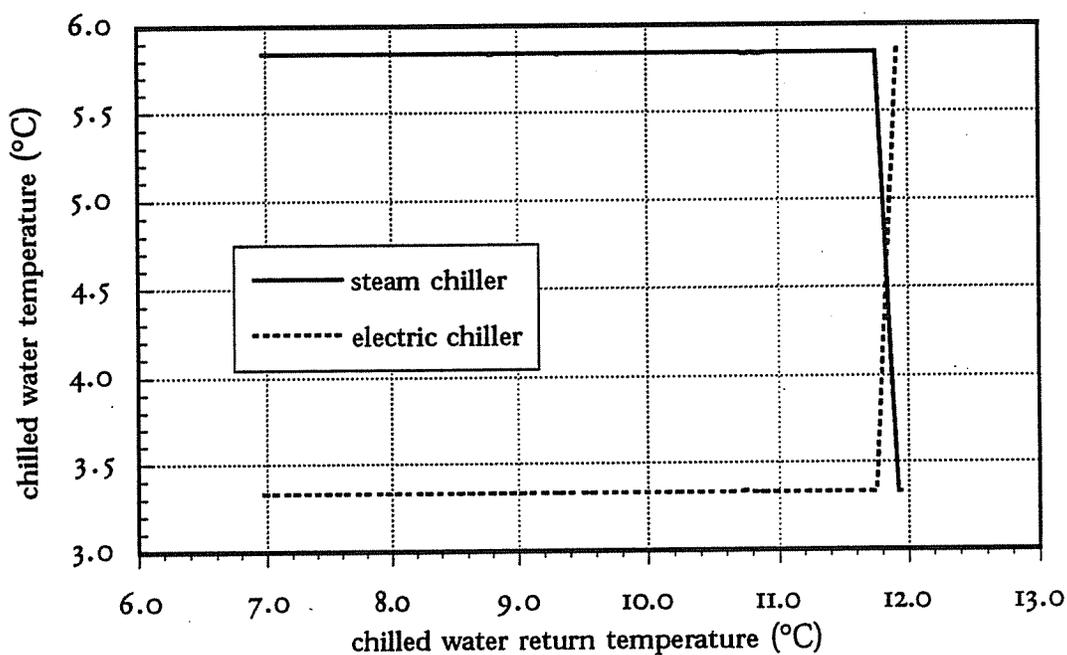


Figure 3: Optimal chilled water temperatures from both chillers as a function of chilled water return temperature. Electric costs are 2.64¢/kW·hr.

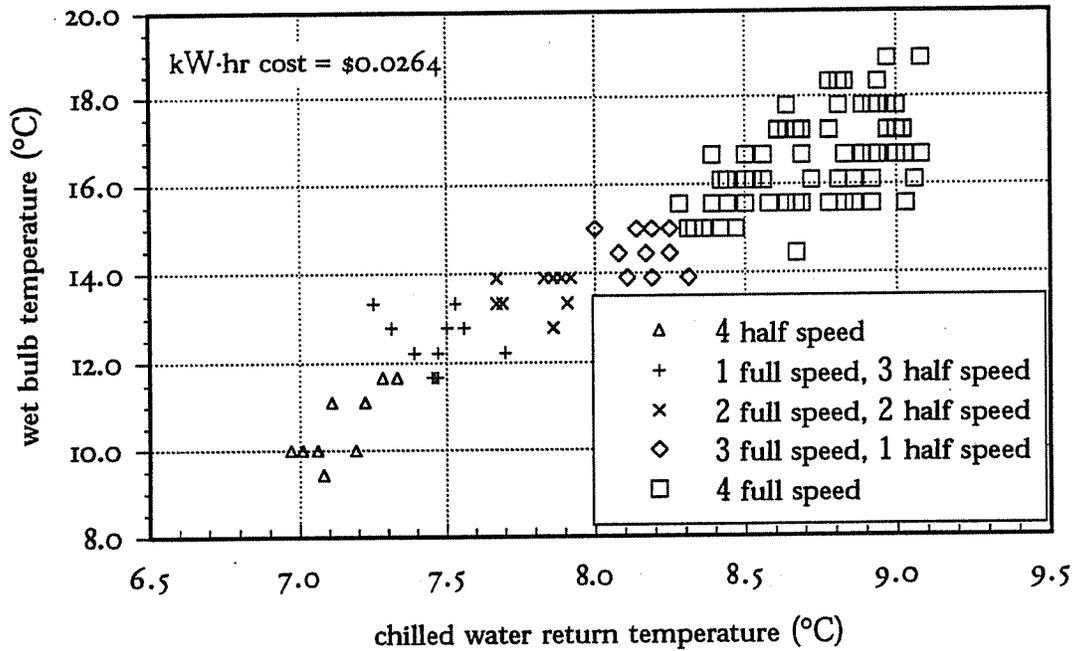


Figure 4: Optimal cooling tower fan speed control vs. chilled water return temperature and wet bulb temperature. Electrical cost is \$0.0264/kW·hr.

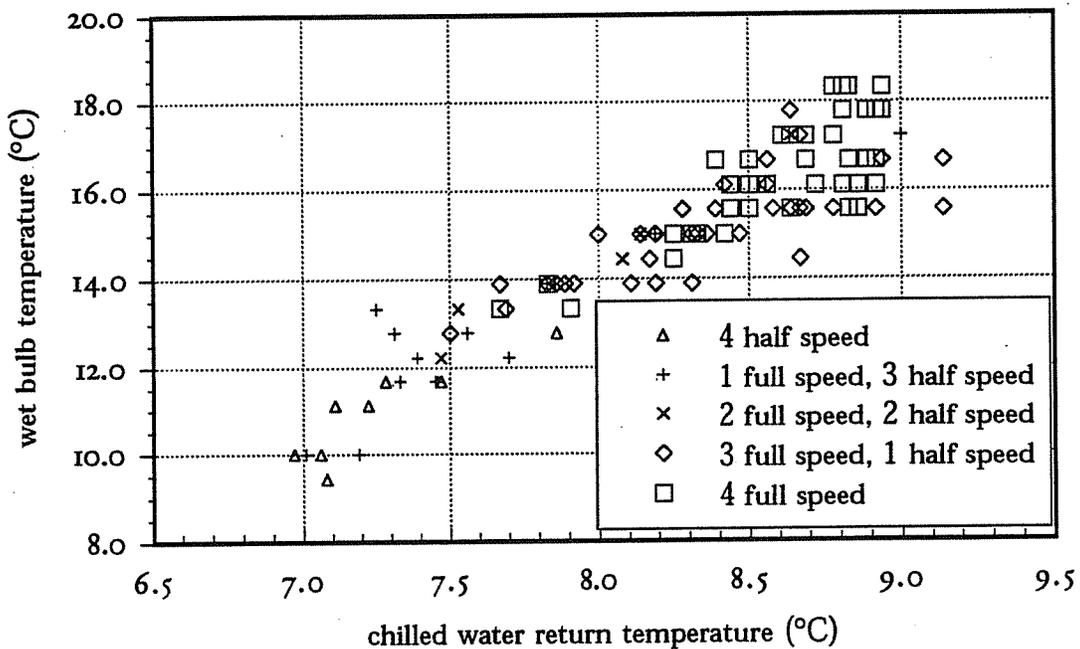


Figure 5: Conventional cooling tower fan speed control vs. chilled water return temperature and wet bulb temperature.

sured data of weather and loads over a month for the Walnut Street Plant, what would have been the optimal control can be determined after the fact.

For hours with relatively small chilled water loads, the optimal control is not influenced by peak demand charges. The optimum control at lower loads is the same as the results from solving the optimization problem without demand charges. To simulate the presence of demand charges, the chiller plant model includes an additional cost at each hour for any power increment above a user set "demand limit" parameter. This demand limit parameter could be set to the highest demand previously calculated during the period or an expected peak demand for the period. At relatively large loads, this limit will influence the optimum control solution for that hour.

Total power use for each hour optimized without demand charges is plotted in Figure 6. The hours have been rank ordered by hourly power use. The optimum peak demand which minimizes costs may be approximated by progressively lowering the demand limit in successive optimization runs until the incremental savings in reduced demand costs is less than the additional energy (steam and kW·hr) costs. At relatively high demand charges (compared to kW·hr cost), the savings in lowering the peak will always be greater than the increased energy costs and the optimum will be at the lower bound on demand. By rank ordering the hours by power demand, only the first few of n hours need to be minimized with the changing p_{\max} constraint. Optimum control values at hours having an optimum power demand less than p_{\max} are not affected by the constraint and need not be recomputed.

The lower bound on demand is found by minimizing electrical use each hour rather than energy cost. The largest minimum power requirement over the month for the Walnut Street Plant is 5,062 kW. Rerunning the optimization including a demand cost for power exceeding 5,062 kW influences the optimal control during eight of the highest load hours. The savings in demand charges compared to optimizing without regard to demand is \$4,086. The increase in energy charges during those high load hours is \$8. For the Walnut Street plant, the optimum peak demand during the month is at its lower bound.

Comparing optimal costs and conventional costs including demand charges, the savings is \$8,406 for the month, averaging \$271 per day. This represents a 4.3% decrease in the monthly energy bill. The percentage savings is not very large, but because yearly energy costs are in the millions of dollars,

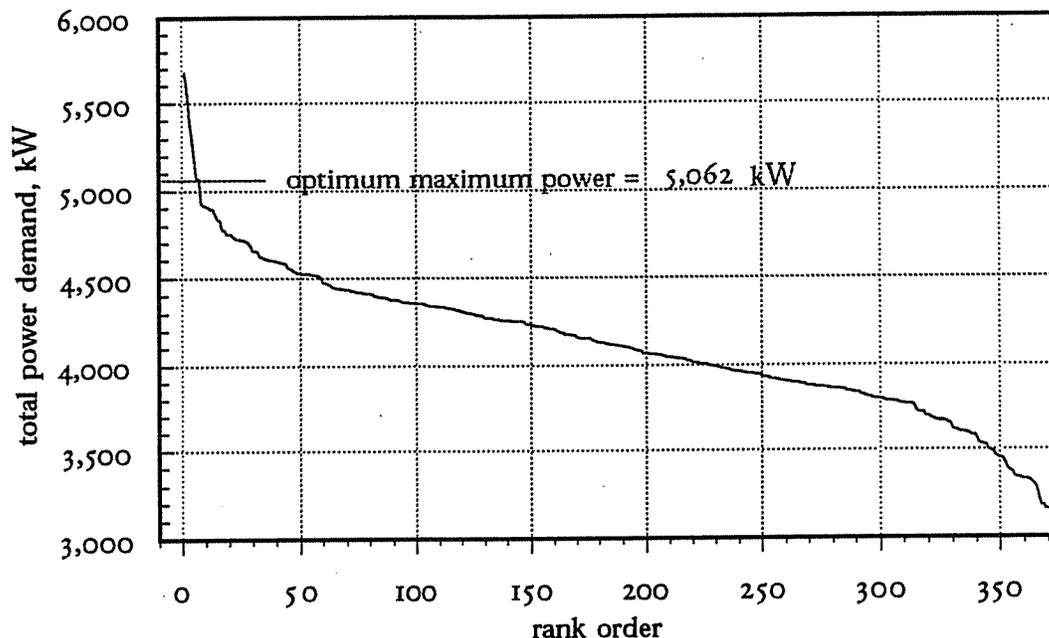


Figure 6: Optimal hourly power demand without regard to demand charges is rank ordered. Including demand charges in the optimization limits the peak demand to 5,062 kW.

the absolute savings is substantial. The small percentage in cost savings is primarily due to the small range of possible individual chiller supply temperatures. The required mixed water temperature of 4.6 °C is only 1.3 °C greater than the minimum allowable temperature, leaving approximately a 2.6 °C range of operation. A larger range for possible individual chiller supply temperatures would allow for increased savings. Optimal control allowing a warmer chilled water supply temperature is discussed below.

In the design of facility air conditioning systems, a chilled water temperature of 7.2 °C (45.0 °F) is commonly assumed to be available for meeting the largest cooling load. For purposes of comparison, conventional and optimal control were recalculated with the plant required to supply 7.2 °C instead of 4.6 °C chilled water. The return water temperature was set to be 2.4 °C (4.6 °F) higher so that the chilled water load would be approximately the same as indicated in the measured data. With the warmer mixed chilled water supply temperature the individual chillers can operate over a wider range of temperatures.

Under conventional control at the warmer supply temperature, steam and electrical energy costs are \$60,450 and \$91,050 respectively. Adding the power demand charge of \$34,741 results in a total

cost of \$186,231. Optimal control reduces the total cost to \$160,002 if demand costs are not considered in the optimization. Limiting the power demand (to the largest minimum required hourly demand) reduces demand costs by \$4,116 and increases steam and kW·hr costs by \$258. The total monthly cost of \$156,144 is a 16.2% savings over conventional control. The larger percentage savings of optimal over conventional control is possible with the warmer supply temperature because the individual chillers can operate over a wider range of supply temperature. Costs for conventional and optimal control under various conditions are summarized in Table 1.

Table 1: Comparison of costs under different operating strategies

		savings over conventional control	total cost, \$	kW demand cost	subtotal, kW·hr and steam costs	kW·hr cost	steam cost
Chilled water supply temperature = 4.6 °F							
conventional control		-	196,554	38,074	158,478	95,034	63,444
optimal control	without demand constraint	4,333	192,221	39,523	152,700	99,930	52,770
	with demand constraint	8,407	188,147	35,437	152,708	99,818	52,890
Chilled water supply temperature = 7.2 °F							
conventional control		-	186,231	34,731	151,500	91,050	60,450
optimal control	without demand constraint	26,229	160,002	36,432	123,570	93,186	30,384
	with demand constraint	30,087	156,144	32,316	123,828	92,706	31,122

Optimal supply temperatures at the higher electric kW·hr rate (\$0.044/kW·hr) are plotted in Figure 7. As return water temperature increases, optimum operation switches from loading the steam chiller to loading the electric chiller at about 11 °C, then back to the steam chiller at temperatures greater than about 13 °C.

The small jump in load above 11 °C T_{chwr} is due to the minimum operating load required before a chiller can be operated. In the transition region near 11 °C T_{chwr} the optimum operation switches from favoring one chiller to the other. Two data points in this range have high wet bulb temperatures compared to other points having approximately the same chilled water return temperatures. At high wet bulb temperature, the cooling water temperature entering the refrigerant and steam condensers rises. Because of the steam condenser, steam driven chiller performance is relatively more sensitive to changes in wet bulb temperature than the electric chiller. At the two points near 11 °C having high

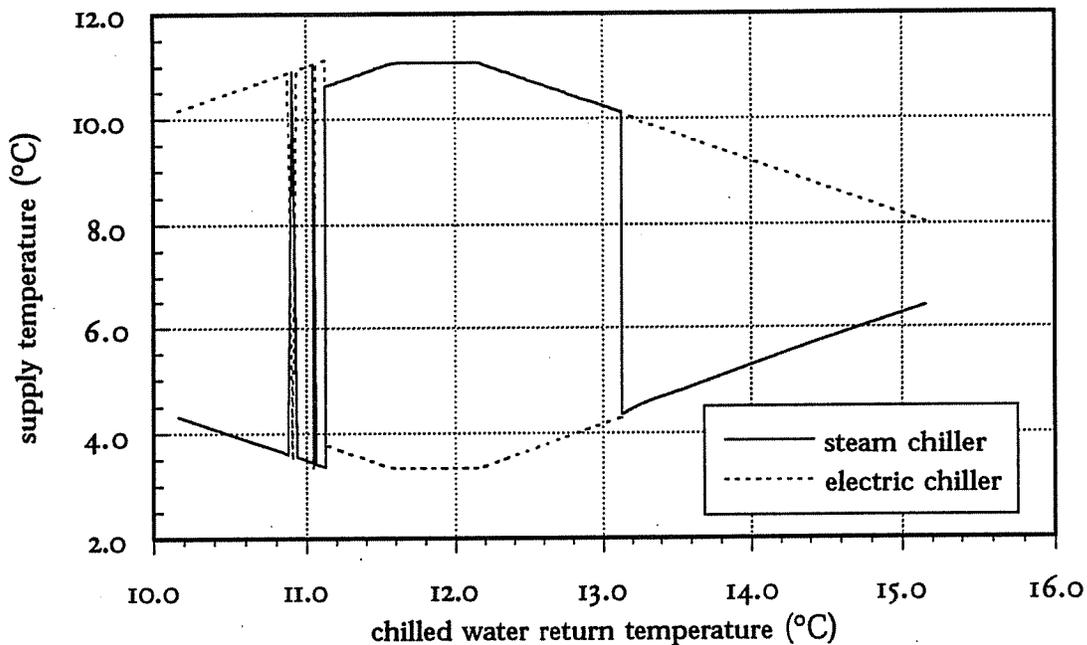


Figure 7: Optimal chilled water supply temperatures of the steam and electric chiller versus chilled water return temperature. The mixed chilled water supply temperature is fixed to be 7.2 °C.

coincident wet bulb temperature, meeting the load with the electric chiller is more economical than meeting it with the steam driven chiller.

5. Summary and Conclusions

A mechanistic model can have advantages over other models (polynomial equations, neural nets, etc.) in extrapolating model predictions. The mechanistic model, in general, constrains the relationships between inputs and outputs to adhere to conservation laws and energy transport relationships observed in physical phenomena. Thus the mechanistic type model for the chilled water plant is advantageous over interpolative type models when the optimal plant operation occurs at control values not present in the measured data.

In determining the optimum supervisory control for the Walnut Street Plant, the model extrapolated responses as the minimization algorithm searched for optimal values of the control variables. The optimal control is to fully load one chiller (operate it at the minimum allowable chilled water supply temperature) and use the other chiller to meet any remaining load. This optimal control usually caused both chillers to operate at a chilled water supply temperature a few degrees outside the

range of measured data used in the fit. Although the accuracy demonstrated for the interpolated predictions is not guaranteed for the extrapolated values, using the mechanistic models gives some confidence that the variable values conform to physical laws. With arbitrary curve fits or neural nets, any prediction outside the region of the data used in the fit is highly suspect.

The optimal control strategy that was found is to operate the chillers in a "priority" mode where one chiller operates at its maximum possible load (priority loaded) and the other chiller meets any remaining load. Which chiller should be given priority is a function of the chilled water return temperature. For high time of day electrical rates, the electric chiller should be given priority control at low chilled water return temperatures. At higher loads (higher chilled water return temperatures) the steam chiller should operate at maximum load. For lower electrical rates, the switch between electric and steam chiller priority occurs at a higher chilled water return temperature.

The optimal control strategy of fully loading one of the chillers rather than running both at part load is attributed to the concave shape of the chiller part load curve. For the chiller system modeled in this work, individual chiller performance (COP) improves with increasing load and the maximum efficiency occurs at full load. The typical part load curve for conventional centrifugal chillers is concave, with a maximum COP at about 70 percent load. Previous guidance toward optimally controlling multiple chillers with this characteristic has been to operate all chillers at the same chilled water temperature. [Bra88] Maintaining equal chilled water supply setpoints coincides with optimal operation for identical chillers having the typical part load performance. However, given two identical chillers having a concave part load curve (maximum efficiency at full load), controlling both at identical chilled water temperatures is the least efficient operation.

The concave part load performance curve is not unique to the chillers studied in this work. Other investigations have found similar performance for operational chillers. [Brn98] One conclusion that should be carried forward from this work is, that for multiple chiller systems, equivalent chilled water setpoints does not always result in near optimal operation. The part load performance of individual chillers should be considered before implementing conventional control strategies.

References

- [Bra88] J. Braun. *Methodologies for the Design and Control of Central Cooling Plants*. Ph.D. Thesis. University of Wisconsin-Madison. 1988.

- [Brn98] M. Brandemuehl. Implementation of on-line supervisory control of cooling plants without storage. ASHRAE Research Project 823-RP. Interim report presented at the ASHRAE Annual Meeting on 18 January 1998 in San Francisco.
- [Fla97] B Flake, J. Mitchell and W. Beckman. Parameter estimation for multiresponse ononlinear chilled-water plant models. *ASHRAE Transactions*. 99(1), 1997.
- [Gof94] W. Goffe, G. Ferrier and J. Rogers. Global optimization of statistical functions with simulated annealing. *Journal of Econometrics*. Vol. 60. 1994.
- [Hac85] R. Hackner, J. Mitchell and W. Beckman. HVAC System Dynamics And Energy Use In Buildings—part II. *ASHRAE Transactions*. 91(1), 1985.