

**Model Predictive Control of HVAC Systems:
Implementation and Testing at the University of California, Merced**

Phillip Haves (LBNL), Brandon Hancey (UTRC), Francesco Borrelli (UC Berkeley),
John Elliot (UC Merced), Yudong Ma (UC Berkeley), Brian Coffey (LBNL),
Sorin Bengea (UTRC), Michael Wetter (LBNL)

June 29, 2010

DISCLAIMER

This document was prepared as an account of work sponsored by the United States Government. While this document is believed to contain correct information, neither the United States Government nor any agency thereof, nor The Regents of the University of California, nor any of their employees, makes any warranty, express or implied, or assumes any legal responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represents that its use would not infringe privately owned rights. Reference herein to any specific commercial product, process, or service by its trade name, trademark, manufacturer, or otherwise, does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States Government or any agency thereof, or The Regents of the University of California. The views and opinions of authors expressed herein do not necessarily state or reflect those of the United States Government or any agency thereof or The Regents of the University of California.

This report was also prepared as a result of work sponsored by the California Energy Commission (Energy Commission) and the University of California (UC). It does not necessarily represent the views of the Energy Commission, UC, their employees, or the State of California. The Energy Commission, the State of California, its employees, and UC make no warranty, express or implied, and assume no legal liability for the information in this report; nor does any party represent that the use of this information will not infringe upon privately owned rights. This report has not been approved or disapproved by the Energy Commission or UC, nor has the Energy Commission or UC passed upon the accuracy or adequacy of the information in this report.

Acknowledgments

This work was supported by the Assistant Secretary for Energy Efficiency and Renewable Energy, Office of Building Technology, State and Community Programs of the U.S. Department of Energy under Contract No. DE-AC02-05CH11231 and by the California Energy Commission PIER Buildings program through the California Institute for Energy and the Environment.

Executive Summary

A Model Predictive Control algorithm was developed for the UC Merced campus chilled water plant. Model predictive control (MPC) is an advanced control technology that has proven successful in the chemical process industry and other industries [1-3]. The main goal of the research was to demonstrate the practical and commercial viability of MPC for optimization of building energy systems. The control algorithms were developed and implemented in MATLAB, allowing for rapid development, performance, and robustness assessment.

The UC Merced chilled water plant includes three water-cooled chillers and a two million gallon chilled water storage tank. The tank is charged during the night to minimize on-peak electricity consumption and take advantage of the lower ambient wet bulb temperature. The control algorithms determined the optimal chilled water plant operation including chilled water supply (CHWS) temperature set-point, condenser water supply (CWS) temperature set-point and the charging start and stop times to minimize a cost function that includes energy consumption and peak electrical demand over a 3-day prediction horizon.

A detailed model of the chilled water plant and simplified models of the buildings served by the plant were developed using the equation-based modeling language Modelica. Steady state models of the chillers, cooling towers and pumps were developed, based on manufacturers' performance data, and calibrated using measured data collected and archived by the control system. A detailed dynamic model of the chilled water storage tank was also developed and calibrated. Simple, semi-empirical models were developed to predict the temperature and flow rate of the chilled water returning to the plant from the buildings. These models were then combined and simplified for use in a model predictive control algorithm that determines the optimal chiller start and stop times and set-points for the condenser water temperature and the chilled water supply temperature. The report describes the development and testing of the algorithm and evaluates the resulting performance, concluding with a discussion of next steps in further research.

The experimental results show a small improvement in COP over the baseline policy but it is difficult to draw any strong conclusions about the energy savings potential for MPC with this system only four days of suitable experimental data were obtained once correct operation of the MPC system had been achieved. These data show an improvement in COP of $3.1\% \pm 2.2\%$ relative to a baseline established immediately prior to the period when the MPC was run in its final form. This baseline includes control policy improvements that the plant operators learned by observing the earlier implementations of MPC, including increasing the temperature of the water supplied to the chiller condensers from the cooling towers.

The process of data collection and model development, necessary for any MPC project, resulted in the team uncovering various problems with the chilled water system. Although it is difficult to quantify the energy savings resulting from these problems being remedied, they were likely on the same order as the energy savings from the MPC itself. Although the types of problems

uncovered and the level of energy savings may differ significantly from other projects, some of the benefits of detecting and diagnosing problems are expected from the use of MPC for any chilled water plant.

The degree of chiller loading was found to be a key factor for efficiency. It is more efficient to operate the chillers at or near full load. In order to maximize the chiller load, one would maximize the temperature difference across chillers and the chilled water flow rate through the chillers. Thus, the CHWS set-point and the chilled water flow-rate can be used to limit the chiller loading to prevent chiller surging. Since the flow rate has an upper bound and the CHWS set point has a lower bound, the chiller loading is constrained and often determined by the chilled water return temperature (CHWR). The CHWR temperature is primarily comprised of warm water from the top of the TES tank. The CHWR temperature falls substantially as the thermocline approaches the top of the tank, which reduces the chiller loading. As a result, it has been determined that overcharging the TES tank can be detrimental to the chilled water plant efficiency. The resulting MPC policy differs from the current practice of fully charging the TES tank. A heuristic rule could possibly avoid this problem without using predictive control. Similarly, the COP improvements from the change in CWS set-point were largely captured by a static set-point change by the operators. Further research is required to determine how much of the MPC savings could be garnered through simplified rules (based on the MPC study), with and without prediction.

The UC Merced campus is currently only a fraction of its future planned size and the TES tank is sized for a substantially larger campus cooling load. Consequently, an upcoming extension of the research is to further characterize via simulations the advantages of MPC scheme to chilled water systems similar to UC Merced with different storage capacity, cooling loads and climate.

An implementation approach has been developed that is less intensive computationally than the approach used for the initial implementation and testing described in the rest of this report. The goal was to produce an implementation of MPC that could be used on a routine basis by the building operators without assistance from the research team and could also be used to reduce the computational effort required for the assessment of the performance and benefits of MPC in different scenarios, such as the building out of the campus. The approach will be tested in the summer of 2010, taking advantage of the repairs and enhancements to the instrumentation of the chilled water plant and distribution system implemented by the UC Merced Facilities staff in the first half of 2010. As discussed in the report, problems with chilled water flow rate measurements, in particular, limited both the performance and the assessment of the MPC algorithm in 2009.

Contents

1	INTRODUCTION	11
1.1	PROJECT DESCRIPTION, GOALS AND OBJECTIVES	11
1.1.1	<i>Goal and Objectives</i>	12
1.1.2	<i>Project Context</i>	12
1.1.3	<i>Technical Challenges</i>	13
1.2	BACKGROUND	14
1.2.1	<i>Model Predictive Control</i>	14
1.2.2	<i>Modelica</i>	15
1.3	PROJECT OVERVIEW	15
2	DETAILED MODEL OF THE CHILLED WATER PLANT.....	18
2.1	INTRODUCTION.....	18
2.2	ELECTRICALLY DRIVEN CENTRIFUGAL CHILLER	20
2.2.1	<i>Model</i>	20
2.2.2	<i>Model Calibration</i>	22
2.3	PRIMARY PUMP	26
2.3.1	<i>Model</i>	26
2.3.2	<i>Model Calibration</i>	27
2.4	SECONDARY PUMP	29
2.4.1	<i>Model</i>	29
2.4.2	<i>Model Calibration and Validation</i>	29
2.5	THERMAL ENERGY STORAGE TANK	31
2.5.1	<i>Model</i>	31
2.5.2	<i>Model Validation</i>	33
2.6	COOLING TOWER.....	34
2.6.1	<i>Model</i>	34
2.6.2	<i>Model Validation</i>	35
2.7	CAMPUS LOAD MODEL	37
2.7.1	<i>Model Description</i>	37
	GainsConverter component.....	38
	BuildingThermalLoad component	39
	Overall inputs, outputs and parameters	40
2.7.2	<i>Model Calibration</i>	40
2.8	MODELING THE CAMPUS CHILLED WATER RETURN TEMPERATURE.....	42
2.8.1	<i>Model Description</i>	42
	Model Principles	43
	Variation of UA as a function of flow rates and load.....	45
	An Empirical Addition to the Model to Account for the Reduction in UA at Low Loads	45
	Temperature Oscillations at Building Heat Exchangers at Low Loads	46
2.8.2	<i>Model Calibration</i>	49
2.9	INTERCONNECTION OF DETAILED MODELS	49
2.10	BASELINE POLICY MODEL.....	51
2.10.1	<i>Weather Forecasts</i>	52
2.10.2	<i>Energy Pricing</i>	52
3	MODEL PREDICTIVE CONTROL DEVELOPMENT	54
3.1	MPC DETAILS	55
3.2	MOVE BLOCKING STRATEGY	56
3.3	TERMINAL CONSTRAINTS.....	57
4	EXPERIMENTS.....	60
4.1	EXPERIMENT DESCRIPTION AND RATIONALE	61
4.2	EXPERIMENTAL PROCEDURES	63

4.2.1	<i>Execution of Experiment II</i>	63
4.2.2	<i>Performance Criteria</i>	64
4.3	EXPERIMENT II DATA ANALYSIS AND MPC PERFORMANCE RESULTS.....	65
5	DISCUSSION	73
5.1	CHILLED WATER PLANT OPERATIONAL PROBLEMS IDENTIFIED	73
5.1.1	<i>Primary Loop Flow-Rate Sensor</i>	73
5.1.2	<i>Secondary and Tertiary Loop Flow-Rate Sensors</i>	74
5.1.3	<i>Chiller Surging Prediction</i>	77
5.1.4	<i>Ambient Humidity Sensor</i>	78
5.1.5	<i>Degraded Cooling Tower Performance</i>	79
5.1.6	<i>Chilled Water Storage Tank Operation</i>	81
5.2	APPLICABILITY OF RESULTS.....	81
5.3	CHALLENGES FOR COMMERCIAL FEASIBILITY	81
6	CONCLUSIONS AND RECOMMENDATIONS	83
6.1	FUTURE WORK.....	84
7	ACKNOWLEDGEMENTS	85
8	REFERENCES	86
9	APPENDICES	88
9.1	APPENDIX 1: IDENTIFICATION OF A BASELINE PERFORMANCE MODEL	88
9.2	APPENDIX 2: MPC EXPERIMENT I.....	91
9.2.1	<i>Execution of Experiment I</i>	91
9.2.1.1	MPC-Based Set-points Generation and Implementation Procedure.....	91
9.2.1.2	Baseline Week 1	92
9.2.1.3	Model Predictive Policy (Experiment Week 2).....	93
9.2.1.4	Baseline Week 3	94
9.2.2	<i>Model Validation with Experiment Data</i>	95
9.2.2.1	Validity of Models	95
9.2.3	<i>Performance Estimation</i>	99
9.2.3.1	Performance Estimation for Baseline and Experiment Weeks.....	100
9.2.3.2	Performance Comparison and Limits to Savings	103
9.3	APPENDIX 3: CHILLED WATER SUPPLY TEMPERATURE EXPERIMENT.....	106
9.3.1	<i>Motivations for Experiment</i>	106
9.3.2	<i>Design of Experiment</i>	109
9.3.3	<i>Execution of Experiment</i>	110
9.3.4	<i>Experimental Results and Analysis</i>	110
9.3.4.1	Results: Campus Level	110
9.3.4.2	Results: Individual buildings	114
9.3.4.3	Analysis: Model modifications to fit data before second MPC experiment.....	122

Figures

FIGURE 1.1: CHILLED WATER PLANT WITH MODEL PREDICTIVE CONTROL AND PARAMETER ESTIMATOR	11
FIGURE 2.1: DIAGRAM OF CHILLED WATER PLANT USING THE BASELINE POLICY	18
FIGURE 2.2: CHILLER PIPING LABELS	20
FIGURE 2.3: CHILLER FUNCTION BLOCK DIAGRAM	20
FIGURE 2.4: CHILLER 1 MAXIMUM CAPACITY FRACTION	23
FIGURE 2.5: CHILLER 1 ENERGY INPUT RATIO AS A FUNCTION OF TEMPERATURE	23
FIGURE 2.6: CHILLER 2 MAXIMUM CAPACITY FRACTION	23
FIGURE 2.7: CHILLER 2 ENERGY INPUT RATIO AS A FUNCTION OF TEMPERATURE	23
FIGURE 2.8: CHILLER 2 ENERGY INPUT RATIO AS A FUNCTION OF PLR (BLUE DOTS MANUFACTURER DATA, DASHED RED LINE MODEL PREDICTION)	24
FIGURE 2.9: CHILLER 2 COEFFICIENT OF PERFORMANCE AS A FUNCTION OF PLR (BLUE DOTS MANUFACTURER DATA, DASHED RED LINE MODEL PREDICTION)	24
FIGURE 2.10: CHILLER #2 COMPARISON OF PREDICTED CATALOG POWER TO MEASURED CHILLER POWER	25
FIGURE 2.11: CHILLER #3 COMPARISON OF PREDICTED CATALOG POWER TO MEASURED CHILLER POWER	25
FIGURE 2.12: COMPARISON OF MEASURED CHILLER #2 PERFORMANCE TO CATALOG MANUFACTURER PERFORMANCE	25
FIGURE 2.13: COMPARISON OF MEASURED CHILLER #3 PERFORMANCE TO CATALOG MANUFACTURER PERFORMANCE	25
FIGURE 2.14: CHILLER PUMP 1 CURVES AFTER TRIMMING	28
FIGURE 2.15: PUMP MODEL FIT BASED ON CALIBRATION POINTS	28
FIGURE 2.16: SECONDARY PUMP POWER MODEL VALIDATION	30
FIGURE 2.17: APPROXIMATE SECONDARY PUMP POWER UPPERBOUND	30
FIGURE 2.18: TANK FLOW DIAGRAM	32
FIGURE 2.19: ILLUSTRATION OF FINITE ELEMENT MODEL VERSUS MOVING BOUNDARY MODEL	32
FIGURE 2.20: THERMAL ENERGY STORAGE TANK VALIDATION	34
FIGURE 2.21: COOLING TOWER DIAGRAM	35
FIGURE 2.22: COOLING TOWER FUNCTIONAL BLOCK DIAGRAM	35
FIGURE 2.23: MEASURED VERSUS MODEL PREDICTED LEAVING WATER TEMPERATURE	36
FIGURE 2.24: VALIDATION OF CUBIC RELATIONSHIP BETWEEN FAN SPEED AND POWER	36
FIGURE 2.25: FIVE DAY TIME PLOT OF MEASURED VERSUS PREDICTED APPROACH TEMPERATURE	37
FIGURE 2.29: MODEL CONTEXT	42
FIGURE 2.30: MEASURED RETURN TEMP, AUG23-SEP4 2009 FOR THE CLASSROOM AND OFFICE BUILDING AND FOR THE ENTIRE CAMPUS	44
FIGURE 2.31: UA VS WATER AND AIR FLOW RATES (LEFT), AND RETURN TEMP VS LOAD (RIGHT)	45
FIGURE 2.32: UA VS WATER AND AIR FLOW RATES (LEFT), AND RETURN TEMP VS LOAD (RIGHT), AFTER EMPIRICAL UA REDUCTION FACTOR ADDED TO MODEL	46
FIGURE 2.34: ABSOLUTE ERROR VS LOAD	49
FIGURE 2.31: SAMPLE DRY BULB TEMPERATURE AND RELATIVE HUMIDITY FORECAST	52
FIGURE 3.2 CAMPUS LOAD	58
FIGURE 3.3 CAMPUS LOAD LOWER BOUND OF THE CONTROLLED PERIODIC INVARIANT SET	59
FIGURE 4.1 SIGNAL FLOW FROM MPC ALGORITHM TO CHILLER PLANT LOCAL CONTROLLERS	61
FIGURE 4.2: ILLUSTRATION OF ELECTRICAL AND THERMAL POWER FLOWS FOR EFFICIENCY DEFINITIONS FOR CHILLER, STORAGE TANK AND CAMPUS	62
FIGURE 4.3 SEQUENCE OF CHARGING-WINDOW ACTIVITIES DURING MPC EXPERIMENT II WEEK	63
FIGURE 4.4. CONDENSER LEAVING WATER TEMPERATURE OPTIMIZED SET-POINTS (SOLID, BLUE)	66
FIGURE 4.5. CHILLED WATER SUPPLY TEMPERATURE SET-POINT (SOLID, BLUE) AND ACTUAL VALUES (DASHED, BLACK)	66
FIGURE 4.6. CHILLED WATER SUPPLY MASS FLOW RATE SET-POINT (SOLID, BLUE) AND ACTUAL VALUES (DASHED, BLACK)	67
FIGURE 4.7: CONDENSER WATER TEMPERATURE SET-POINTS AND MEASURED DATA, OCTOBER 6-10 CHARGING PERIODS	67
FIGURE 4.8: DETAILS OF MPC OPERATION, CHARGING PERIOD ON THE NIGHT OF OCTOBER 6, 2009	68
FIGURE 4.9: DETAILS OF MPC OPERATION, CHARGING PERIOD ON THE NIGHT OF OCTOBER 8, 2009	69
FIGURE 4.10: DETAILS OF MPC OPERATION, CHARGING PERIOD ON THE NIGHT OF OCTOBER 9, 2009	70
FIGURE 4.11: DETAILS OF MPC OPERATION, CHARGING PERIOD ON THE NIGHT OF OCTOBER 10, 2009	71

FIGURE 4.12: AVERAGE CHARGING PERIOD COP, MEASURED VS. BASELINES	72
FIGURE 5.1: POWER AND HEAT FLOW DIAGRAM FOR CHILLER	74
FIGURE 5.2: CHILLER #3 COMPARISON OF ESTIMATED COOLING LOAD VERSUS MEASURED COOLING LOAD	74
FIGURE 5.3: CHILLER #2 COMPARISON OF ESTIMATED COOLING LOAD VERSUS MEASURED COOLING LOAD	74
FIGURE 5.4: COMPARISON OF RETURN TEMPERATURES MEASURED CENTRALLY, IN THE TANK, AND THE RETURN TEMPERATURES FROM THE BUILDINGS WEIGHTED BY THEIR MEASURED FLOW RATES	75
FIGURE 5.5: COMPARISON OF THE DAILY CHILLER LOADS, THE DAILY LOADS CALCULATED AT EACH BUILDING AND DAILY CAMPUS LOAD CALCULATED CENTRALLY	76
FIGURE 5.6: COMPARISON OF MEASURED AND ESTIMATED TANK LEVELS	76
FIGURE 5.7: CHILLER #2 SURGING CAPTURED BY $PLRADJ > 1$, START DAY IS AUG. 20, 2008; (BLUE CHW = CHWR.CH2; GREEN CHW = CHWS.CH2; BLUE CW = CWR.CH2; GREEN CW = CWS.CH2)	78
FIGURE 5.8: COMPARISON OF CAMPUS WEATHER STATION TO AIRPORT WEATHER STATION MEASUREMENTS	79
FIGURE 5.9: COMPARISON OF PREDICTED AND MODELED CONDENSER WATER SUPPLY TEMPERATURE USING MODEL CALIBRATED FROM UC MERCED RELATIVE HUMIDITY SENSOR.	79
FIGURE 5.10: COOLING TOWER PERFORMANCE CURVES PROVIDED BY BALTIMORE AIRCOIL COMPANY	80
FIGURE 5.11: COMPARISON OF MEASURED AND PREDICTED LEAVING TEMPERATURE	80
FIGURE 9.1. COP LINEAR REGRESSION VALIDATION RESULTS	89
FIGURE 9.2. ABSOLUTE ERROR OF COP LINEAR REGRESSION	90
FIGURE 9.4. DATA FLOW FOR CHILLER PLANT MODEL VALIDATION	95
FIGURE 9.5. PLOTS OF ACTUAL VALUES OF THE CONTROL VARIABLES DURING THE	96
FIGURE 6. PLOTS OF WET BULB TEMPERATURE AND TANK CAPACITY DURING THE	97
FIGURE 7. PLOTS OF THE AMBIENT TEMPERATURE, CHILLER WATER RETURN TEMPERATURE AND ACTUAL COOLING LOAD DURING MPC POLICY IMPLEMENTATION WEEK OF EXPERIMENT 1	97
FIGURE 9.8 PLOTS OF TOTAL ELECTRICAL AND THERMAL POWERS, AND COEFFICIENT OF PERFORMANCE DURING THE MPC POLICY IMPLEMENTATION WEEK OF EXPERIMENT 1	98
FIGURE 9.9. PLOTS OF ACTUAL AND SIMULATED TOTAL ELECTRICAL AND THERMAL ENERGIES, AND CUMULATIVE COP FOR THE EXPERIMENT 1 MPC POLICY WEEK	99
FIGURE 9.10. PLOTS OF SET-POINTS VALUES FOR CHILLER WATER SUPPLY TEMPERATURE AND FLOW RATE FOR EXPERIMENT 1 MPC AND BASELINE WEEK 2	100
FIGURE 9.11. PLOTS OF AMBIENT TEMPERATURE, TANK CAPACITY, CHILLER WATER RETURN TEMPERATURE, AND COOLING LOADS FOR EXPERIMENT 1 MPC AND BASELINE WEEK 2	101
FIGURE 9.13. PLOTS OF GENERATED AND CONSUMED ENERGY, AND CUMULATIVE COP FOR EXPERIMENT 1 MPC POLICY AND BASELINE 2 WEEKS	102
FIGURE 9.14. PERFORMANCE COMPARISON BETWEEN BASELINE AND MPC POLICY WEEKS. REMARKS: (I) THE COLOR CODE IS: BLUE FOR EXPERIMENT WEEKS, AND GREEN FOR VARIOUS OTHER WEEKS; (II) THE RELATIVE SIZE CIRCLES REPRESENTS THE RELATIVE SIZE OF THE COP MAGNITUDES	104
FIGURE 9.15: RETURN TEMPERATURE VS LOAD	106
FIGURE 9.16: MAINTAINING A LOG MEAN TEMPERATURE DIFFERENCE	107
FIGURE 9.17: EFFECT OF INCREASING T_{wi} AT LOAD = 2.5MW	108
FIGURE 9.18: EFFECT OF INCREASING T_{wi} AT LOAD = 0.5MW	108
FIGURE 9.19: EFFECT OF INCREASING T_{wi} AT LOAD = 1.0MW	109
FIGURE 9.15. CAMPUS SUPPLY TEMPERATURE (PRE-EXPERIMENT IN BLUE, EXPERIMENT WEEK IN RED)	111
FIGURE 9.16. RETURN CHILLED WATER TEMPERATURES TO CENTRAL PLANT (PRE-EXPERIMENT IN BLUE, EXPERIMENT WEEK IN RED)	111
FIGURE 9.17. RETURN TEMPS TO CENTRAL PLANT (MODEL PREDICTION IN LIGHT BLUE, EXPERIMENT WEEK IN RED)	113
FIGURE 9.18. TOTAL CAMPUS FLOW RATES (PRE-EXPERIMENT IN BLUE, EXPERIMENT WEEK IN RED)	113
FIGURE 9.19. CLASSROOM AND OFFICE BUILDING: RETURN TEMPS (PRE-EXPERIMENT IN BLUE, EXPERIMENT WEEK IN RED)	114
FIGURE 9.20. CLASSROOM AND OFFICE BUILDING: CHILLED WATER FLOW RATES	115
FIGURE 9.21. LIBRARY BUILDING: RETURN CHILLED WATER TEMPERATURES	115
FIGURE 9.22. RECREATION BUILDING: RETURN CHILLED WATER TEMPERATURES	116
FIGURE 9.23. SIERRA TERRACES: RETURN TEMPERATURES	117
FIGURE 9.24. KINGS DORM: RETURN TEMPERATURES	117
FIGURE 9.25. KINGS DORM: SUPPLY TEMPERATURES AFTER RECIRCULATION	118

FIGURE 9.26. SCIENCE AND ENGINEERING BUILDING: RETURN TEMPERATURES	119
FIGURE 9.27. SCIENCE AND ENGINEERING BUILDING: FLOW RATES	119
FIGURE 9.28. DINING BUILDING: RETURN TEMPERATURES	120
FIGURE 9.29. COMMONS BUILDING: RETURN TEMPERATURES	121
FIGURE 9.30. FACILITIES BUILDING: RETURN TEMPERATURES	121
FIGURE 9.31. MEASURED DATA FROM EXPERIMENT	122
FIGURE 9.32. MEASURED OUTDOOR AIR TEMPERATURE VS CAMPUS LOAD	123
FIGURE 9.33. γ SENSITIVITY	124
FIGURE 9.34. β SENSITIVITY	124
FIGURE 9.35. FULLY STRETCHED MODEL (RED) VS THE MODEL AS CONFIGURED BEFORE EXPERIMENT (BLUE)	125
FIGURE 9.37. WATER SIDE DELTA T VS LOAD	126
FIGURE 9.38. UA (CALCULATED) VS LOAD	127
FIGURE 9.39. UA (CALCULATED) VS WATER FLOW RATE	127
FIGURE 9.40. AIR FLOW RATE (CALCULATED) VS WATER FLOW RATE	128
FIGURE 9.41. MEASURED DATA (TOP) VS MODEL OUTPUTS (BOTTOM)	129
FIGURE 9.42. MEASURED VS MODEL FOR STANDARD SUPPLY TEMP (TOP) AND HIGHER TEMP (BOTTOM)	130

1 Introduction

1.1 Project Description, Goals and Objectives

Technologies for energy efficiency improvement in buildings are central to the development of marketable design approaches for net zero energy commercial buildings by 2025 which is a strategic goal of the DOE Buildings Technologies Program. Heating, ventilation, and air conditioning (HVAC) account for 27% of the energy consumption and 45% of peak electrical demand in commercial buildings [4], and 5% of the floor area of US commercial buildings is cooled by central chillers or district chilled water plants [5]. Studies on several central plants indicate that substantial energy savings (up to 60% in some cases where the operation is poor) are possible with properly deployed optimized system level controls and variable speed technology [6].

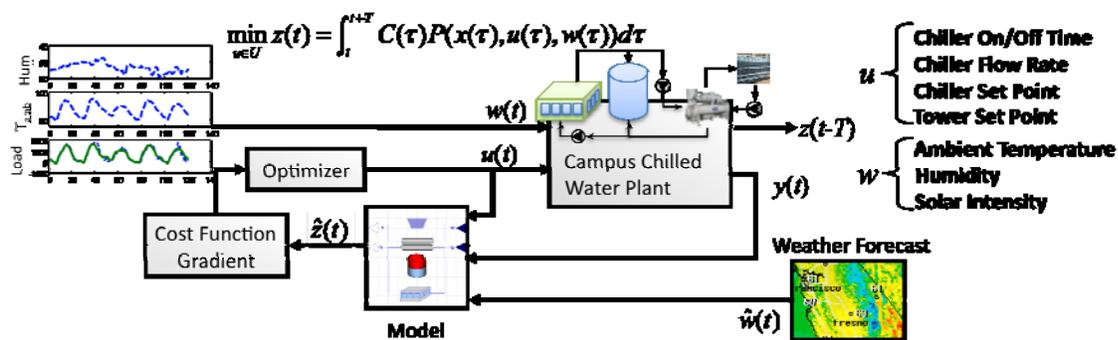


Figure 1.1: Chilled water plant with model predictive control and parameter estimator

If thermal energy storage (TES) is effectively used in conjunction with weather and cooling load forecasting, TES can be used to reduce electricity costs by managing energy consumption and peak loads [7]. Peak demand charges alone account for about 40% of commercial building electricity costs [4]. The utilization of weather and cooling load forecasting enables extreme cooling loads to be shifted by generating and storing chilled water more efficiently when the weather and electricity rates are advantageous. Currently, weather or cooling load forecasting are not commonly used to dynamically manage the chilled water storage. Instead, static, heuristic policies are typically employed.

For example, the TES tank at University of California Merced is completely charged on nights when the charge level is observed to be below 50% full capacity. As a result, the chillers at the UC Merced chilled water plant (CWP) are operated on a nightly/semi-nightly basis to charge the thermal energy storage (TES) tank. Based on work performed during this project, up to 5% energy savings can be achieved by adjusting the CWS temperature set point. Additional savings are possible by leveraging cooling load and weather forecasts to charge the TES tank on the

coolest nights and to optimize temperature profile of the TES tank. In order for current ECMS technology to realize the most of the above potential savings, heuristic control logic that far exceeds typical complexity must be skillfully and laboriously coded, monitored, and maintained.

The current state of control technology within building systems is a collection of logic-based rules and single-input-single-output feedback loops regulating set-points **Error! Reference source not found.** The best industry practices define rules to accommodate for the diverse modes of operation encountered during hourly, daily, weekly, and seasonal operation. Rules can also play a role in coordinating set-points among separate subsystems, such as scheduling the leaving cooling tower set point according to the chiller full load amps percentage. Consequently, the degree of optimality of the resulting ECMS is dependent on the expertise of the control designer, a thorough understanding of the system, persistent monitoring for changes in subsystem behavior and performance, and frequently updating the rules to reflect any changes in the subsystems.

Model predictive control is an advanced control technology that has proven successful in the chemical process industry. It is a promising technology for building energy control and the main goal of the proposed program is to demonstrate the practical and commercial viability of model predictive control algorithms for optimization of building energy systems. The scope of the project entails developing algorithms that reduce total and peak energy consumption and manage thermal energy storage tank utilization in a campus chilled water plant.

1.1.1 Goal and Objectives

The goal of the project was to contribute to the assessment of the potential of advanced control methods to improve the energy performance of buildings by investigating the benefits of applying Model Predictive Control (MPC) to chilled water plants with thermal storage. The project objectives were:

- To implement and test Model Predictive Control (MPC) of the UC Merced chilled water plant and measure the energy savings and peak demand reduction.
- To estimate the energy savings and peak demand reduction potential of HVAC systems with active thermal storage systems in commercial buildings, based on simulation.

1.1.2 Project Context

The work reported here was a collaboration between Lawrence Berkeley National Laboratory (LBNL), United Technologies Research Center (UTRC), and the University of California Berkeley that was funded by the US Department of Energy and the California Energy Commission. It was one of a set of three linked projects to implement, test, and demonstrate enhanced energy performance through improved control and visualization, based on real-time modeling.

The field testing and demonstration activities in all three projects were performed on the campus of the University of California, Merced. Aggressive energy performance and sustainability goals have resulted in two campus buildings and the central plant each achieving LEED Gold ratings. The campus was also designed to be a 'living laboratory' and has a significantly enhanced level of instrumentation in order to support the development and demonstration of energy-efficient technologies and practices. Faculty members with expertise in energy technologies have been selectively recruited to UC Merced and energy-oriented undergraduate and graduate curricula are under development.

The UC Merced chilled water system consists of a chiller plant (three chillers redundantly configured as two in series, one backup in parallel), a two million gallon chilled water tank, a primary distribution system and secondary distribution loops serving each building. The two series chillers are operated each night to recharge the storage tank, which meets campus cooling demand the following day. Because thermal energy is stored actively within the water tank and passively within the building fabric, minimizing energy consumption is a dynamic optimal control problem, i.e. the optimal value of a set-point at any point in time depends on the future thermodynamic state trajectories, which in turn depend on the future values of disturbances such as coil loads and weather.

1.1.3 Technical Challenges

The main technical challenge being addressed by this work is the synthesis of robust optimal supervisory control of HVAC systems with thermal storage. Modulation of operational set-points and modification of sequences of operation (equipment start/stop logic for example) is fundamentally a dynamic optimization problem whose solution depends on the future evolution of a nonlinear and large scale system. The key enablers are the use of computationally efficient, equation-based models that can be used in conjunction with nonlinear programming algorithms [9,10], thereby allowing the application of optimal control theory for large scale systems [11], and enterprise-scale web-enabled networked control systems. Beyond this demonstration, model-based methods and computational tools must be developed for design and analysis of robust optimal control algorithms, reducing the impact of uncertainty in models and disturbances and enabling the use of optimal control on a routine basis.

1.2 Background

Chilled water plants serving large commercial buildings or groups of buildings typically consist of several centrifugal, vapor compression chillers, several open cycle, evaporative cooling towers, and circulation pumps. Various chilled water piping schemes are possible, with the *primary/secondary* configuration, in which the primary loop through the chillers is coupled to a secondary loop serving the air handling unit cooling coils, is most common in existing buildings. *Primary/secondary/tertiary* schemes are commonly used for campuses, with the secondary loop circulating water between the central plant and the buildings and the tertiary loops circulating water between the secondary loop and the air handling unit cooling coils. Chilled water plants may also incorporate thermal storage in the form of a tank that either contains ice or chilled water, the purpose being to shift load from the period of high demand in the afternoons to the night time, when utility costs are lower and/or the lower ambient wet bulb temperature allows more efficient chiller operation.

One common operational problem in chilled water plants is ‘low delta-T’, a relatively small difference between the temperature of the chilled water supplied by the plant and the temperature of the chilled water returning to the plant. This may be caused by poor design and/or poor operation, including undersized cooling coils and excessive chilled water flow rates. More generally, the power consumption of a chilled water plant is the sum of the chiller power, cooling tower fan power and circulation pump power, and minimization of this power typically involves the use of variable speed fans and pumps, and possibly variable speed chillers, and an operational strategy that produces an optimal, or near-optimal, trade-off between the various speeds. Near-optimal operation may be obtained through the implementation of carefully designed ‘reset’ strategies for the chilled water and condenser water temperature set-points and heuristic ‘dispatch’ strategies that specify the combination of chillers and cooling tower cells to be used at each operating point. Formal optimal control may be achieved by the use of an algorithm that minimizes a cost function, typically utility cost, while satisfying constraints designed to ensure adequate comfort in the occupied spaces. The use of thermal storage systems turns the static optimization problem into a dynamic operation problem in which the objective is to minimize a cost function over time.

1.2.1 Model Predictive Control

Model Predictive Control (MPC) is one of several methods to compute an optimal control solution. MPC employs a real-time dynamic simulation model to predict critical thermodynamic states and system power consumption as a function of forecasted disturbances and set-points. The model is used to evaluate a cost function over the prediction horizon and an optimizer is used to compute set-point trajectories to minimize the cost. A variety of computational methods exist to compute the optimal control including sensitivity methods, adjoint-based methods and dynamic programming, each of which may take advantage of a particular structure of the problem and model. The cost function used in this demonstration is a weighted

sum of peak electrical demand and energy consumption over the time horizon; the weights can be adjusted so that the cost function reflects the actual utility cost. Constraints are included in the model and optimization to ensure that the set-point trajectories respect zonal comfort (especially humidity) and equipment constraints. MPC has not been widely applied in buildings to date; a simulation-based assessment of MPC applied to buildings is presented in [12].

1.2.2 Modelica

Modelica is an equation-based modeling language that is well positioned to become the de-facto open standard for modeling multi-physics dynamical systems. It already has a successful track record in many industrial-scale multi-physics applications (see also <http://www.modelica.org/publications>). The previous experience of the team in using this language showed that it allows the analysis of innovative building systems and the dynamic behavior of controls that are outside the scope of existing building simulation programs [13, 14]. Such a modeling paradigm requires that each component only sees the boundary conditions that a real object would interact with (such as a temperature and heat flux boundary conditions or a data bus). This makes the integration of system models for different physical domains possible on the mathematical modeling layer (by linking model equations). The implications are that developing new component models and integrating these models into a system model can be done much more rapidly compared to the prevailing modeling approach that is used by the buildings industry, which is largely based on procedural FORTRAN code.

The Modelica language was used to rapidly develop equation-based models. Dymola [15] was utilized as a compiler for the Modelica language, which allowed the model development to focus on the mathematical description of the physical phenomena rather than the solution of simultaneous equations. This approach provided increased flexibility in integrating models for various components of the chilled water plant. For control purposes, the Modelica models were either recoded in MATLAB [16] and C or sampled to provide look up tables for the rapid calculations required by the MPC algorithm.

1.3 Project Overview

The control algorithms were developed and implemented in MATLAB, allowing for rapid development, performance, and robustness assessment. Detailed dynamic models of the chilled water piping system and a simplified dynamic model of the buildings at UC Merced were developed in the Modelica language, and used as the basis of MATLAB models or lookup tables for use within the controller. The control algorithms determine the optimal chilled water plant set-points including leaving chilled water temperature (LWT), tower return temperature and the charging start and stop times to minimize a cost function that includes energy consumption and peak electrical demand over a 3-day prediction horizon.

In comparison, the existing controls at UC Merced maintain a 39°F chilled water supply

temperature, and the chillers are sequenced manually in a manner that is heuristically optimized for full load. Although the storage tank enables load shifting to off-peak hours to reduce peak demand, the lack of optimized automatic control results in conservative over-charging of the tank, which results in a lower return temperature due to the significant thickness of the thermocline, which then results in reduced efficiency.

Two weeks of MPC experiments were carried out to demonstrate and test the real-time implementation of the MPC controller at the Merced campus. The first week was in early June 2009. This week allowed the research team to debug the models, the controller, and the overall process. It produced a number of positive side effects, such as the identification of a chiller surging problem and a broken flow rate sensor in one of the campus buildings that was causing increased pump energy at that building and keeping the campus return temperature lower than it should have been, thus making the system COP lower than it should have been. However, because of errors in the MPC implementation, this experiment week did not produce any energy savings when compared with the facility's standard control practices. The system operators, however, did learn from this implementation that their condenser water temperature set-point should be increased, and thus changed their standard practice to produce higher system COPs. In September 2009, a separate short experiment was carried out on campus to empirically test the impacts of changing the supply water temperature, the results of which was used to validate and better calibrate the model representing the effect of the building loads on the chilled water system. The second week of MPC experimental implementation was in early October 2009. With re-calibrated models and adjustments in the MPC configuration, this second week proved much more successful than the first, even with its significantly cooler weather.

The following is an overview of the research process that was carried out, the details of which are covered in the following three sections.

- Model development
 - System component models were developed in Modelica and Matlab.
 - These models were calibrated and verified using measured data from the campus control system.
- MPC development
 - The models were integrated into an MPC controller in MATLAB. Lookup tables derived from off-line component simulations were used for some of the components to produce a faster on-line model within the controller. Available MATLAB toolboxes were used to solve the optimization problem.
 - Scripts were used to obtain local weather forecasts and to pull data from the campus control system as needed to update system states in the online model. The optimized set-points were communicated verbally to the operators. This was done in order to allow the operators to build up confidence in the MPC. Automated communication between the MPC system and the campus control system should be implemented and tested as part of any follow-on work.
- Experiments

- The first MPC experimental implementation: This week allowed the research team to debug the models, the controller, and the overall process.
- A short experiment was performed to test the impacts of increasing the temperature of the chilled water supplied to the campus.
- The second experimental implementation: This week demonstrated significant energy savings through the use of MPC.

A discussion of the identification of the baseline performance models is presented in Appendix 1

2 Detailed Model of the Chilled Water Plant

2.1 Introduction

This section describes the main components used to generate and store thermal energy at the UC Merced campus — shown in Figure 2.1. The energy management and control system (EMCS) provides the supervisory control that coordinates two electric chillers, cooling towers, and multiple pumps in the overnight charging of a 7,600 cubic meter (2 million gallon) chilled water storage tank. The thermal energy storage allows the chilled water plant to take advantage of nighttime electricity rates and lower ambient temperature when charging the thermal energy storage, and the chilled water is pumped from the tank and distributed throughout the campus the following day.

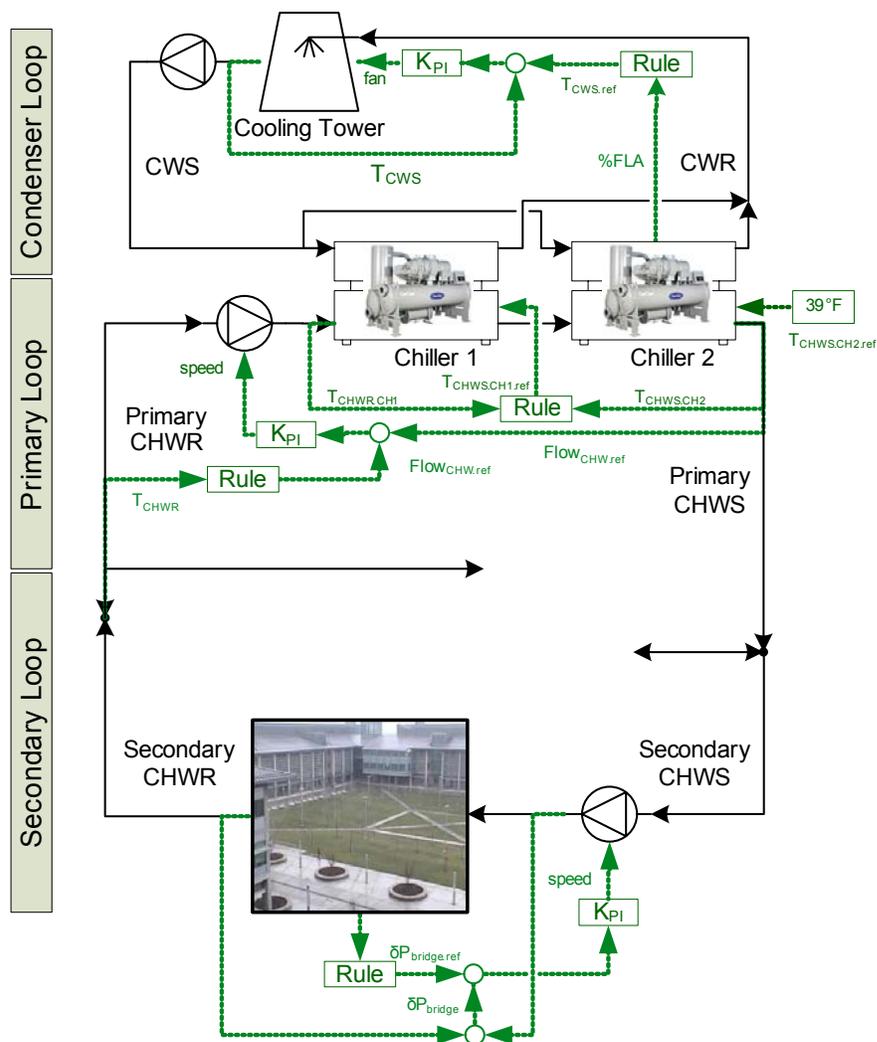


Figure 2.1: Diagram of chilled water plant using the baseline policy

The chilled water system consists of a condenser loop, a primary loop, a secondary (campus) loop, and several tertiary (building) loops. The chilled water is generated via chillers and cooling towers within the primary and condenser loops. The chilled water is stored in a stratified thermal energy storage tank, and distributed to the buildings throughout campus via the secondary loop. The tertiary loop uses pumps and valves within each building to distribute the chilled water for consumption by the fan coils and air handling units (AHUs). The chilled water is warmed by the air-side cooling load of the buildings and returned to the secondary loop.

The thermal energy storage tank holds approximately 7,600 cubic meters (2 million gallon) of chilled water. The storage chilled water can easily meet the current campus cooling loads for more than a day. Thus, any decision made with respect to quantity and temperature of the chilled water stored in the storage tank affects the performance of the entire system over a relatively long time horizon. In contrast, the chilled water plant components—such as the chillers, cooling towers, and pumps—have very short time constants in comparison. In fact, the 15 minute sampling interval of the energy management control system generally results in undersampling of the transients in these components. As a result, the model presented in proceeding discussion focuses on the dominant dynamics of the thermal energy storage tank and the campus cooling load generated by the buildings, and the operation of the much shorter time constant system components is treated as quasi-static.

Lower-level controllers modulate the operation of the chillers and cooling towers in order to achieve a desired condenser water supply temperature produced by cooling towers, $T_{cws,ref}$, mass flow rate of chilled water through the chillers, $\dot{m}_{chw,ref}$, and chilled water supply temperature, $T_{chws,ref}$. The dynamics of the chillers and cooling towers are neglected and it is assumed that there is no tracking error between the controlled variables and their set-points, i.e. $T_{cws} = T_{cws,ref}$, $\dot{m}_{chw} = \dot{m}_{chw,ref}$, and $T_{chws} = T_{chws,ref}$.

What follows is a mathematical description of the individual components comprising the UC Merced chilled water plant, shown in Figure 2.1. The objective is to develop a simplified yet transparent model that can be used for real time optimization in a MPC scheme.

In order to represent the energy flow associated with fluid flow in either direction, the following function is defined:

$$\dot{H} = \text{semilinea}(\dot{m}, h_1, h_2) = \begin{cases} \dot{m}h_1 & \text{if } \dot{m} \geq 0 \\ \dot{m}h_2 & \text{if } \dot{m} < 0 \end{cases} \quad (2.1)$$

The models use a constant pressure approximation for the enthalpy of the water in the chilled water and condenser loops, i.e. $h = C_p T$. Consequently, internal energy is treated as independent of pressure, i.e. $U = mh$.

notation	description
h	specific enthalpy [J/kg]
\dot{H}	enthalpy flow rate [W]
\dot{m}	mass flow rate [kg/s]
p	pressure [Pa]
P	power [W]
\dot{Q}	heat flow rate [W]
U	internal energy [J]
T	Temperature [K]
ρ	fluid density [kg/m ³]

2.2 Electrically Driven Centrifugal Chiller

The UC Merced chilled water plant rotates among three chillers, using two chillers at any one time. The evaporators of the two operational chillers are piped in series in order to generate a larger temperature difference between the chilled water return and chilled water supply temperatures. Each of the three chillers is a Carrier 19XR chiller. As a result of the series configuration, the compressors and heat exchangers of each of the chillers were tuned for either mid-range chilled water supply temperatures or the lower chilled water supply temperatures.

2.2.1 Model

The centrifugal chiller model was formulated based on the DOE-2.1 regression-based chiller model detailed in [10]. The chiller model uses polynomial curve fits in order to predict the chiller capacity and coefficient of performance. As shown in Figure 2.3, the polynomial curve fits are a function of entering condenser water temperature, leaving evaporator water temperature, and cooling load.

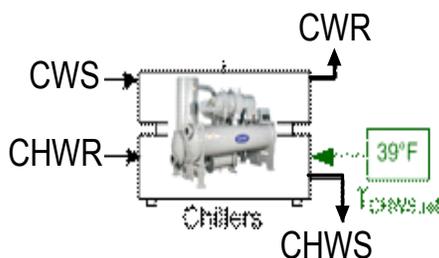


Figure 2.2: Chiller piping labels

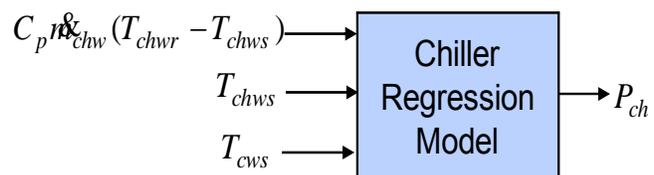


Figure 2.3: Chiller function block diagram

The energy balances across the evaporator and the condenser heat exchangers are respectively represented as

$$0 = \dot{m}_{chw} (h_{chws} - h_{chwr}) - \dot{Q}_{evap} \quad 0 = \dot{m}_{evap} (h_{chws} - h_{chwr}) - Q_{evap} \quad (2.1a, 2.2b)$$

Rather than using an detailed physical model based on the thermodynamic vapor compression cycle, several empirical functions are used to relate the boundary conditions and cooling load of the chiller to the power consumed. The cooling capacity of the chiller is predicted by the biquadratic function

$$\dot{Q}_{cap} = \dot{Q}_0 CapF_T(T_{chws}, T_{cws}), \quad (2.2)$$

where

$$CapF_T(T_{chws}, T_{cws}) = a_0 + a_1 T_{chws} + a_2 T_{chws}^2 + a_3 T_{cwr} + a_4 T_{cwr}^2 + a_5 T_{cwr} T_{chws}. \quad (2.3)$$

The part load ratio (PLR) is calculated as the fraction of the cooling capacity used by the cooling load.

$$PLR = \max(PLR_{min}, \dot{Q}_{evap} / \dot{Q}_{cap}) \quad (2.4)$$

subject to the constraint $0 \leq PLR \leq 1$. The two energy input ratio (EIR) functions, $EIRF_T(T_{chws}, T_{cws})$ and $EIRF_{PLR}(PLR)$, represent the ratios of power consumed versus the cooling load. The first EIR function is a biquadratic function of the form

$$EIRF_T(T_{chws}, T_{cws}) = b_0 + b_1 T_{chws} + b_2 T_{chws}^2 + b_3 T_{cwr} + b_4 T_{cwr}^2 + b_5 T_{cwr} T_{chws} \quad (2.5)$$

and the second EIR function is a cubic function of the form

$$EIRF_{PLR}(PLR) = c_0 + c_1 PLR + c_2 PLR^2 + c_3 PLR^3 \quad (2.6)$$

The chiller power consumption is calculated as

$$P = \dot{Q}_{cap} EIRF_T(T_{chws}, T_{cws}) EIRF_{PLR}(PLR) / COP_0 \quad (2.7)$$

where COP_0 is the nominal coefficient of performance of the chiller.

In terms of control-oriented modeling, the chiller has number of constraints. As mentioned above,

notation	description
$CapF_T$	biquadratic chiller capacity ratio temperature curve [W/W]
COP_0	nominal coefficient of performance [W/W]
$EIRF_T$	biquadratic energy input ratio temperature curve [W/W]
$EIRF_{PLR}$	quadratic energy input ratio PLR curve [W/W]
P	compressor power [W]
PLR	power load ratio [W/W]
CCR	chiller cycling ratio
\mathcal{Q}_0	nominal chiller cooling capacity [W]
Q_{cap}	chiller cooling capacity [W]
\cdot_{evap}	evaporator
\cdot_{cws}	condenser inlet
\cdot_{cwr}	condenser outlet
\cdot_{chwr}	evaporator inlet
\cdot_{chws}	evaporator outlet

2.2.2 Model Calibration

The following chiller calibration follows the procedure outlined in [11]. The temperature-dependent cooling capacity function $CapF_T$ and the temperature-dependent energy input ratio function $EIRF_T$ must be calibrated using data from the chiller at its maximum cooling capacity. However, chillers are not normally operated at their maximum cooling capacity, thus a large excursion from nominal chilled water plant policies would be required to obtain the maximum cooling capacity data. Furthermore, the maximum cooling capacity data is required for a broader range of operating conditions than is experienced under normal operation of the chilled water plant. For example, the chilled water supply set-point for the second chiller is set to 3.9°C (39°F) under typical operations. As a result, it would prove quite difficult to obtain sufficient maximum cooling capacity data for a broad range of operating conditions using only historical data from the UC Merced chilled water plant. Instead, software used by the chiller vendors was used to assess the performance of the chillers according to manufacturer specifications.

A least squares regression fit was used to fit the parameters of $CapF_T$ and $EIRF_T$ to the manufacturer curves. The results of the Chiller #1 model calibration are shown in Figure 2.4 and Figure 2.5. Chiller #2 and Chiller #3 have the same compressor and heat exchanger configurations; thus only a single set of parameters is identified for the $CapF_T$ and $EIRF_T$ curves. The results of the Chiller #2 calibration are shown in Figure 2.6 and Figure 2.7. As shown

in Figure 2.6, Figure 2.7, and Figure 2.8, the regression-based model is able to capture the dominant performance characteristics of Chiller #2. As shown in Figure 2.9, the model is capable of predicting the chiller coefficient of performance reasonably well. As suggested by the model error observed in Figure 2.8, the error could be potentially further reduced by making $EIRF_{PLR}$ a function of T_{CWS} as well [18].

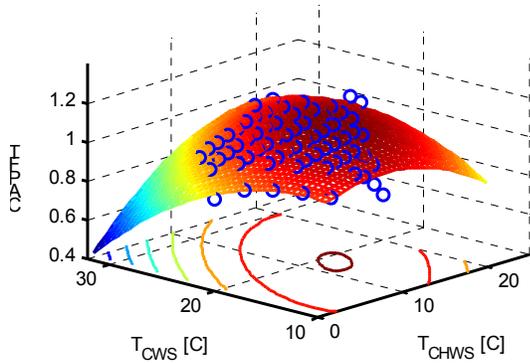


Figure 2.4: Chiller 1 maximum capacity fraction

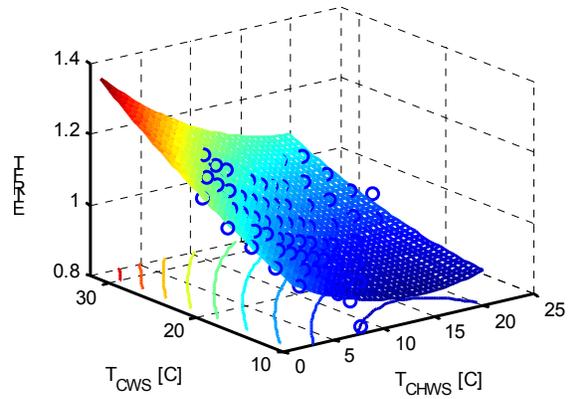


Figure 2.5: Chiller 1 energy input ratio as a function of temperature

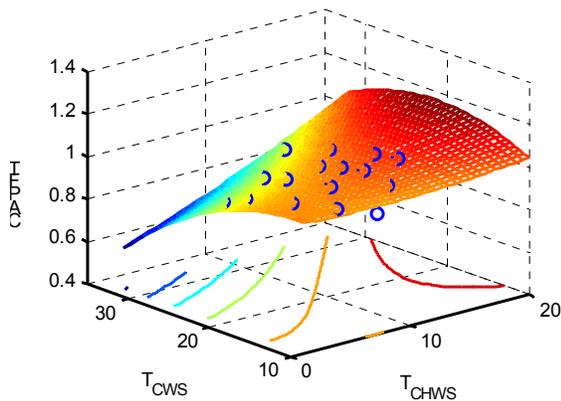


Figure 2.6: Chiller 2 maximum capacity fraction

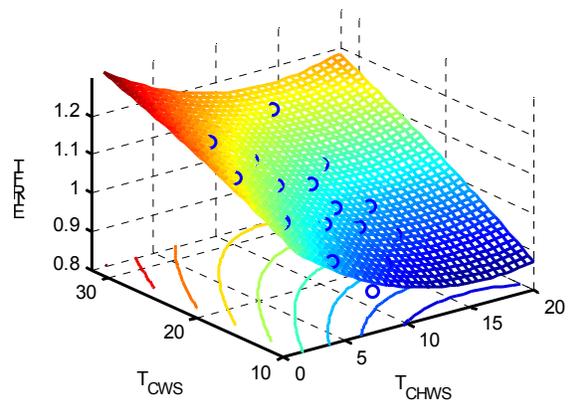


Figure 2.7: Chiller 2 energy input ratio as a function of temperature

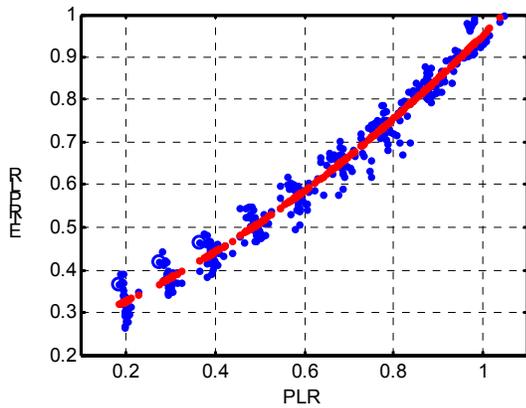


Figure 2.8: Chiller 2 energy input ratio as a function of PLR (blue dots manufacturer data, dashed red line model prediction)

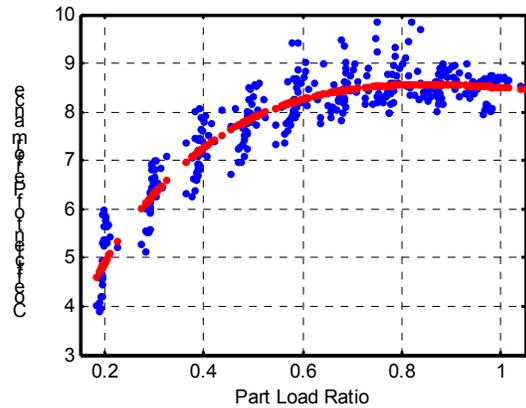


Figure 2.9: Chiller 2 coefficient of performance as a function of PLR (blue dots manufacturer data, dashed red line model prediction)

Figure 2.10 and

Figure 2.11 compare the measured power consumption of Chiller #2 and Chiller #3 to the regression-based models calibrated to the manufacturer curves. Although there were a number of deterioration issues, such as [PH1][bmh2]fouling of the shell and tube heat exchangers, the models calibrated using the manufacturer performance curves predicted power consumption within 10%. The data shown in

Figure 2.10 and

Figure 2.11 were collected for over a one year period. Note there are a number of outliers, which may be attributed to startup transients and other anomalies that were not completely filtered from the original data set. Figure 2.12 and Figure 2.13 show that the catalog data predict the chiller power and COP reasonably well over a 10 day period.

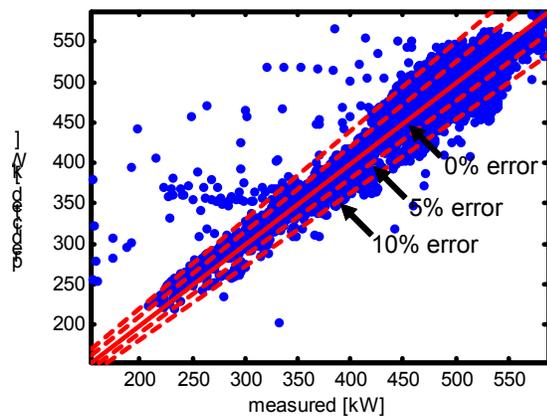
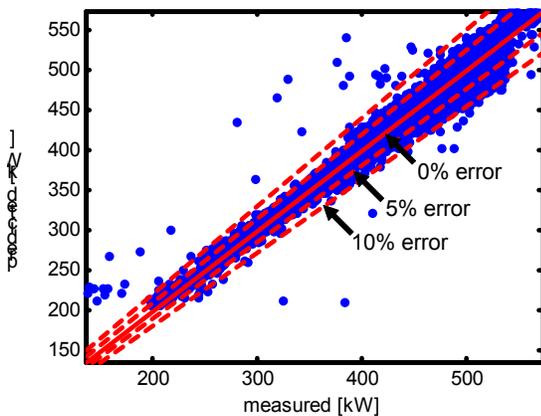


Figure 2.10: Chiller #2 comparison of predicted catalog power to measured chiller power

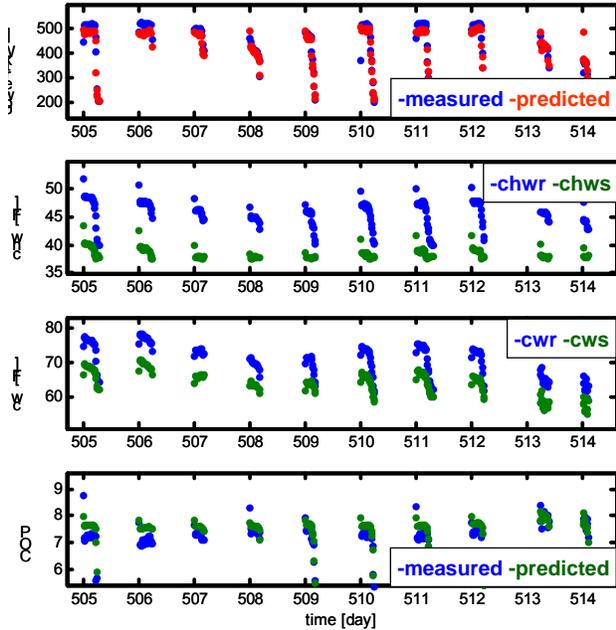


Figure 2.11: Chiller #3 comparison of predicted catalog power to measured chiller power

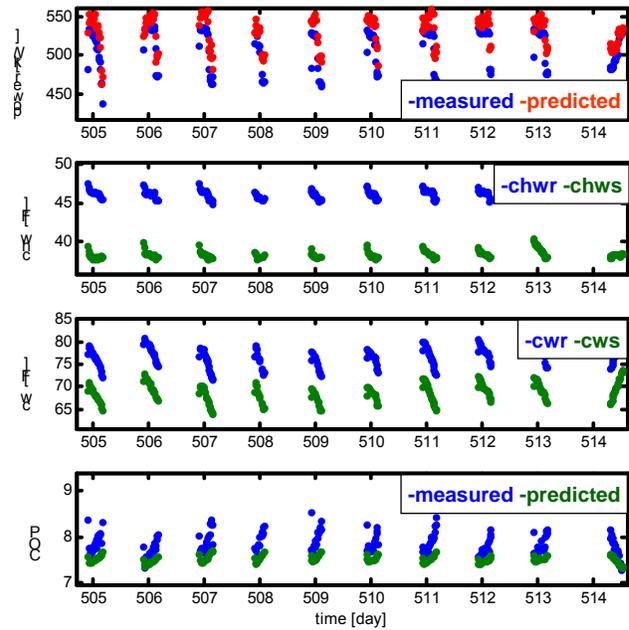


Figure 2.12: Comparison of measured Chiller #2 performance to catalog manufacturer performance

Figure 2.13: Comparison of measured Chiller #3 performance to catalog manufacturer performance

Table 2.1: Calibrated parameter values for Chiller #1

Name	Value	Name	Value	Name	Value
a_0	-7.9872×10^{-1}	b_0	8.5044×10^{-1}	c_1	2.2303×10^{-1}
a_1	4.8241×10^{-3}	b_1	-1.3754×10^{-2}	c_2	6.2412×10^{-1}
a_2	-1.3588×10^{-3}	b_2	6.6912×10^{-4}	c_3	3.1339×10^{-1}
a_3	2.9323×10^{-2}	b_3	7.7426×10^{-3}	Q_0	4.8703×10^6
a_4	-1.3382×10^{-3}	b_4	2.2261×10^{-4}	COP_0	8.1172×10^0
a_5	1.6090×10^{-3}	b_5	-4.7382×10^{-4}		

Table 2.2: Calibrated parameter values for Chiller #2 and Chiller #3

Name	Value	Name	Value	Name	Value
a_0	8.5669×10^{-1}	b_0	9.6335×10^{-1}	c_1	1.9881×10^{-1}
a_1	1.5732×10^{-3}	b_1	-4.4336×10^{-2}	c_2	5.8152×10^{-1}
a_2	-6.4089×10^{-4}	b_2	1.7183×10^{-3}	c_3	2.1191×10^{-1}

a_3	1.1555×10^{-2}	b_3	2.0931×10^{-2}	Q_0	4.8703×10^6
a_4	-6.8540×10^{-4}	b_4	-3.1604×10^{-4}	COP_0	8.1172×10^0
a_5	9.7662×10^{-4}	b_5	1.9195×10^{-4}		

2.3 Primary Pump

The variable speed primary pump determines the flow-rate through the evaporators of the chillers. Consequently, it plays a key role in determining the cooling load placed on the chillers, and the resulting chiller COP.

2.3.1 Model

The pump pressure rise, also known as pump head, for a nominal pump speed N_0 is modeled as a quadratic function

$$\Delta p_0(q_0) = a_0 + a_1 q_0 + a_2 q_0^2, \quad (2.8)$$

where q_0 is the nominal volumetric flow-rate and Δp_0 is the nominal pressure across the pump. Similarly, the efficiency curve of the pump for a nominal pump speed is modeled as a quadratic function

$$\eta(q_0) = b_0 + b_1 q_0 + b_2 q_0^2. \quad (2.9)$$

In order to address variable speed operation of the pump, the following affinity laws are used to scale the pressure and efficiency curves

$$\frac{q}{q_0} = \frac{N}{N_0}, \quad \frac{\Delta p}{\Delta p_0} = \frac{N^2}{N_0^2} \quad (2.10)$$

where the ratio of pump speed N to nominal pump speed is denoted as R . If we take N_0 to be the maximum pump speed, then $0 \leq R \leq 1$. The general pump pressure curve takes the form

$$\Delta p = R^2 \cdot \Delta p_0(q/R), \quad (2.11)$$

where $\Delta p = p_o - p_i$, and the pump power is calculated as

$$P = \frac{\Delta p \cdot q}{\eta(q/R)} \quad (2.12)$$

Note that the efficiency of the variable speed drive (VSD) power electronics is not explicitly parameterized. Generally, the efficiency curve of variable speed drives are known to be relatively flat for normal pump torque and speed operating conditions. Consequently, the efficiency of the VSD has been lumped into the pump efficiency in (2.12). Assuming the variable speed drive and pump motor reject their heat to the ambient mechanical space_[PH3], the energy

balance between the inlet and outlet of the pump is

$$0 = \rho q (h_i - h_o) . \quad (2.13)$$

notation	description
a	coefficients for quadratic pressure difference curve
b	coefficients for quadratic efficiency curve
N	pump motor speed
P	pump power [W]
q	volumetric flow rate [m^3/s]
R	ratio of drive speed to nominal speed [rpm/rpm]
ρ	fluid density [kg/m^3]
\cdot_i	pump inlet
\cdot_o	pump outlet

2.3.2 Model Calibration

Figure 2.14 illustrates the trimmed pump curves for chiller pump 1. The pumps models are based on a quadratic curve relating flow rate to head and efficiency and on the pump affinity laws. Several flow-rates were selected as calibration points for the quadratic curves, and a least squares fit was used to choose the parameters for the quadratic curves. Figure 2.15 illustrates the quality of the curve fit for a select pump. Given that the pump curves for each of the three chiller pumps are similar, separate pump curves were not calibrated.

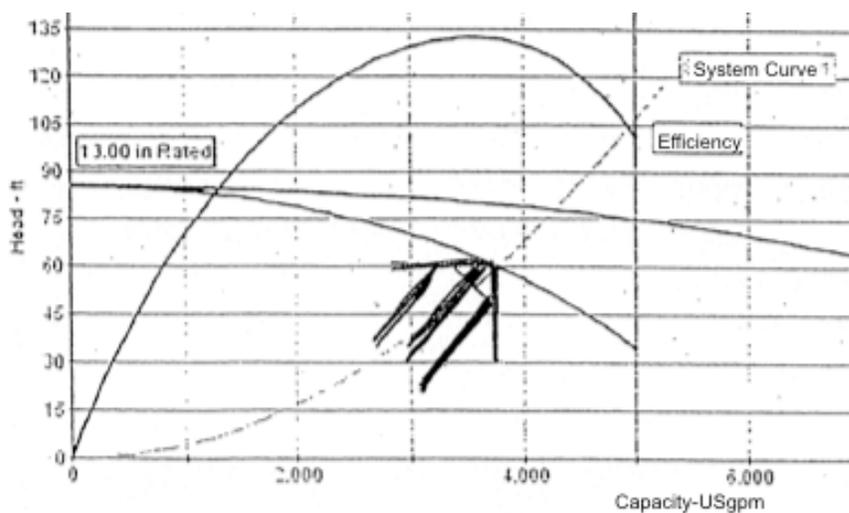


Figure 2.14: Chiller pump 1 curves after trimming

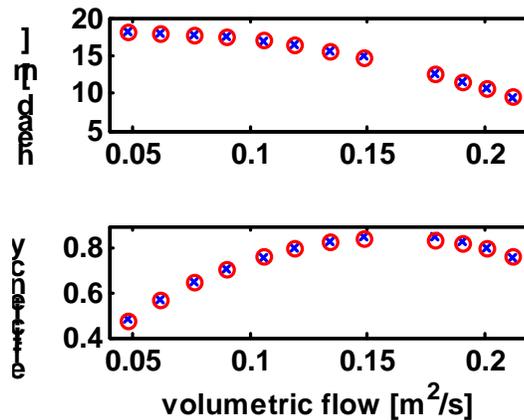


Figure 2.15: Pump model fit based on calibration points

Much like the pump curve, the system curve $\Delta p = c_0 + c_1 q + c_2 q^2$, represents the pressure loss as a function of the flow-rate. The system curve from the design specifications, shown in Figure 2.14, was deemed unreliable, because the primary pumps were unable to achieve the 221 kg/s (3500 gpm) flow-rate. Instead, the primary pumps were only able to achieve 183 kg/s (2900 gpm) before adjusting the pump balancing valve and 196 kg/s (3100 gpm) after adjusting the pump balancing valve. Consequently, system curve was fit using a data set possessing a limited number of points with the pump at less than full speed. The pressure drops across the evaporator of each chiller, specified by the design specifications, were used to calculate the turbulent pressure losses as quadratic terms. Conversely, the pressure loss in the primary loop piping was modeled as a linear function of flow-rate, and the coefficient c_1 was tuned to in order to fit the measured flow-rates data given pump speed.

Table 2.3: Calibrated parameter values for chiller pump 1

Name	Value	Name	Value	Name	Value
a_0	2.5324×10^5 _[PH4]	b_0	8.5797×10^{-2}	c_0	0.0×10^0
a_1	1.1587×10^5	b_1	6.6896×10^0	c_1	6.1792×10^5
a_2	-1.8391×10^6	b_2	-1.4247×10^1	c_2	3.6610×10^6

2.4 Secondary Pump

2.4.1 Model

The secondary (campus) loop pump pressurizes the chilled water for distribution to the tertiary (building) loops. Unlike the primary loop pump, the secondary loop pump speed is modulated to track a differential pressure set-point measured across the secondary loop. In order to obtain the differential pressure set-point, the EMCS employs a type of pressure-reset algorithm. The basic objective is to supply sufficient pressure to outlying buildings to achieve adequate

flow-rates at the buildings, while minimizing over-pressurization which results in increased throttling losses over valves in the tertiary loops. For example, if a tertiary pump in a building is unable to meet the building flow-rate, a cooling request is sent to the secondary loop causing the differential pressure to be incremented. However, there are many buildings and many components (e.g. pumps and valves) inside the tertiary loop contributing to chilled water demands; therefore, it was deemed impractical to construct models detailed enough to predict the cooling requests from each building.

Unlike the primary pump, there was a substantial set of data for a range of flow-rates available for the secondary pump, see Figure 2.16. As a result, historical data was used to calibrate the secondary pump power model. The secondary pump power was fit to the multivariate polynomial

$$P = a_0 + a_1q + a_2q^2 + a_3q^3 + a_4\Delta p \quad (2.14)$$

2.4.2 Model Calibration and Validation

A simple least squares regression was used to fit the parameters a_i . Figure 2.16 illustrates the quality of the secondary pump power prediction given the flow-rate and pressure differential.

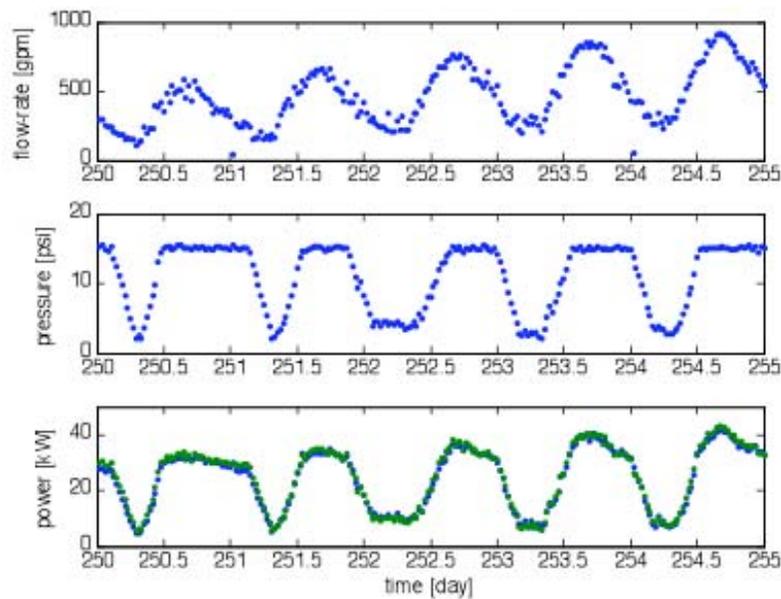


Figure 2.16: Secondary pump power model validation

As aforementioned, the pressure differential was deemed too difficult to predict accurately using first principles. Noting the maximum pressure differential for the pump is limited to 103.4 kPa (15 psi), an approximate upper bound on the power estimate was provided by setting $\Delta p = 103.4$ kPa. As shown in Figure 2.17, the model approximately predicts an upper bound for the power consumed by the secondary pump.

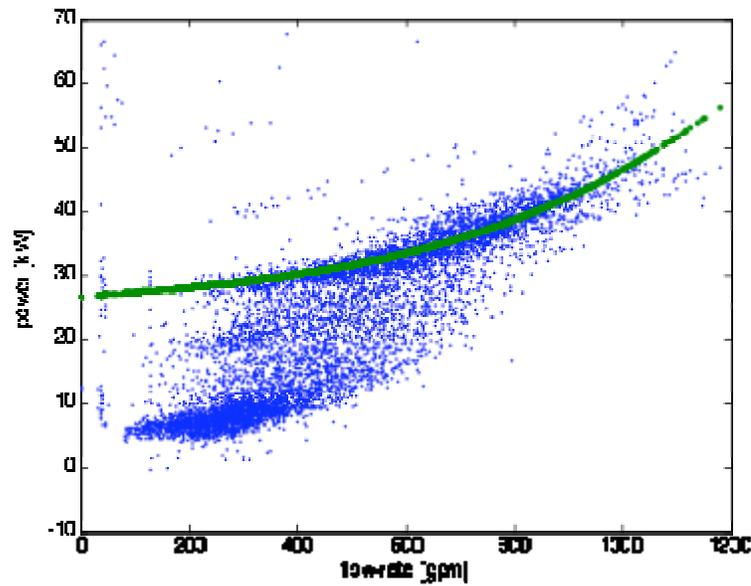


Figure 2.17: Approximate secondary pump power upperbound

Table 2.4: Calibrated secondary pump parameters

Name	Value	Name	Value	Name	Value
a_0	1.4291×10^3	a_1	1.3662×10^1	a_2	1.4764×10^0
a_3	4.2472×10^{-2}	a_4	1.2408×10^0	Δp	$1.034 \times 10^5 \text{ Pa}$

2.5 Thermal Energy Storage Tank

2.5.1 Model

A new TES tank model was developed in the work reported here. During tank charging, cool water from the chillers enters the bottom of the tank and warm water from the top of the tank is returned to the chillers. During tank discharging, cool water from the bottom of the tank is supplied to the campus, and warm from the campus is returned to the top of the tank. The objective of the TES tank model is to capture three essential aspects:

1. The total stored cooling capacity,
2. The temperature of the water supplied to the campus,
3. The temperature of the water returned to the chiller.

These three aspects of interest are effectively embedded in the temperature profile of the stratified tank. In the detailed model (used in the MPC algorithm development and in determining the terminal constraints on the optimization, noted below), the temperature profile in the tank is modeled by discretizing the tank into a number of layers. However, the discretization introduces many dynamic states, which results in a considerable computational

burden for the optimization algorithm. For the MPC online optimization, a low-order model was developed in an effort to capture the primary aspects of the TES tank temperature distribution without using a large number of states. The thermocline between the warm and cool masses of water is treated as a natural moving boundary (treating the steep measured gradient between hot and cold, see Figure 2.19, as a step change in temperature). The cool and warm water are treated as lumped masses, thus requiring only three dynamic states for the tank model, i.e. one for the position of the thermocline and the temperature of each lumped mass. That is, the mass flow rate into (out of) the top port is equal to the mass flow rate out of (into) the bottom port, thereby eliminating one state associated with the mass of water. As a result, the low-order model represents an approximation of the three essential aspects of the TES tank.

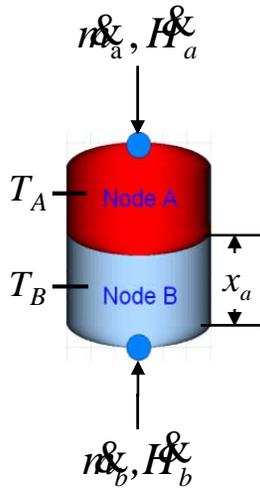


Figure 2.18: Tank flow diagram

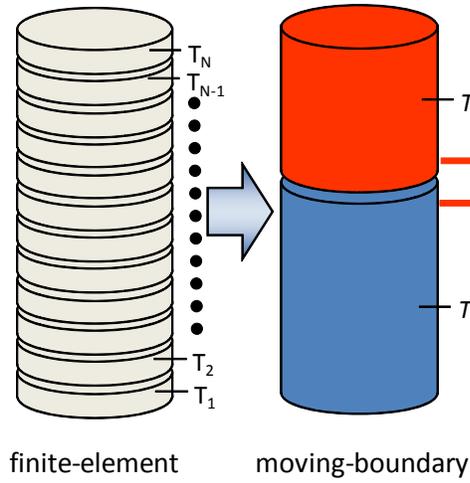
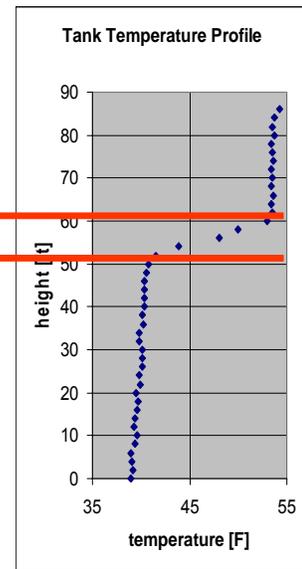


Figure 2.19: Illustration of finite element model versus moving boundary model



The energy balances for the cool and warm masses are

$$\dot{U}_a = \dot{H}_a + \dot{Q}_{amb,a} + \dot{Q}_{b,a}, \quad (2.15)$$

$$\dot{U}_b = \dot{H}_b + \dot{Q}_{amb,b} - \dot{Q}_{b,a}, \quad (2.16)$$

where U_a is the internal energy of the warm water mass, U_b is the internal energy of the cool water mass, $\dot{Q}_{amb,a} = k_1(2\pi R z_a)(T_{amb} - T_a)$ is the heat transfer from the ambient to the warm water mass, $\dot{Q}_{amb,b} = k_1(2\pi R z_b)(T_{amb} - T_b)$ is the heat transfer from the ambient to the cool water mass, and $\dot{Q}_{b,a} = k_2 A_c (T_b - T_a)$ is the heat transfer across the thermocline,

$$\dot{H}_a = \text{semilinear}(\dot{m}_a, h_{chw}, h_a) \quad (2.17)$$

is the enthalpy flow rate into the warm water mass, and

$$\dot{H}_b = \text{semilinear}(\dot{m}_b, h_{chws}, h_b) \quad (2.18)$$

is the enthalpy flow rate into the cool water mass.

Since the University of California Merced chilled water loop is a closed system, $\dot{m}_{tank} = \dot{m}_b = -\dot{m}_a$. Thus the water height in the tank is also constant $z_{tank} = z_a + z_b$. Since the tank has uniform cross-sectional area, the internal energy of the warm and cool masses of water is calculated as $U_a = (\rho A_c z_a) h_a$ and $U_b = (\rho A_c z_b) h_b$, respectively. Last, the mass fraction differential equation is

$$\dot{x}_a = \dot{m}_{tank} / m_{tank}, \quad (2.19)$$

where $m_{tank} = A_c \rho z_{tank}$ is the mass of water in the tank, $x_a = m_{tank,a} / m_{tank} = z_a / z_{tank}$ is the mass fraction of cool water, $x_b = m_{tank,b} / m_{tank} = z_b / z_{tank}$ is the mass fraction of warm water, and $x_a + x_b = 1$.

notation	description
A_c	area of tank cross section
k_1	heat transfer coefficient to ambient [W/K/m ²]
k_2	heat transfer coefficient between cool and warm mass of water [W/K/m ²]
R	radius of cylindrical tank
T_{amb}	ambient temperature [K]
x	mass fraction of water
z	height of water [m]
\cdot_a	warm mass of water above thermocline
\cdot_b	cool mass of water below thermocline

2.5.2 Model Validation

The proposed model is validated by using data collected from May 22nd to May 29th, 2007. The historical inputs were applied to the tank model and the output of the model $[z_a, z_b, T_a, T_b]$ is compared with the measurements.

In Figure 2.20a, the thin solid lines indicate the measurements of the 44 temperature sensors

evenly spaced along the height of the storage tank, and the thick black (red) dotted lines show the temperature of the cool (warm) water. Due to relatively low cooling loads, Figure 2.20 shows the tank was only charged every other night. Around hours 60 and 160 a second thermocline is observed to have formed. The second thermocline was formed by an influx of warm water from the campus after relatively cooler water returned from the campus at night, thereby cooling the warm mass of water above the thermocline.

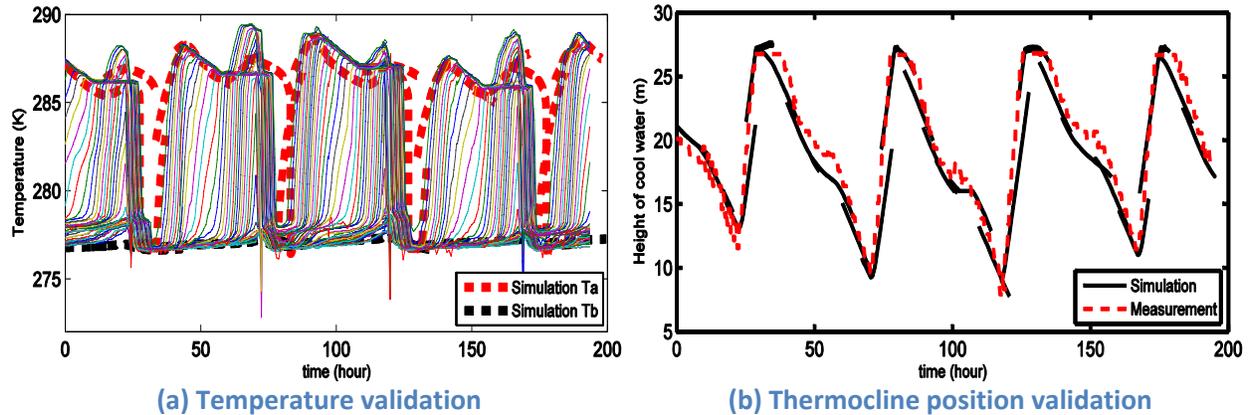


Figure 2.20: Thermal energy storage tank validation

The proposed tank model reflects the temperature dynamics of the top (bottom) layer of the tank water that exits (enters) the tank. However, the model is only capable of capturing the average temperatures of the cool and warm masses of water. As a result, the model does not predict the temperature of the water above the second thermocline. Instead, the model predicts the average temperature of the water above the first and second thermoclines. A higher order model could overcome this limitation, but a higher order model would cause greater computational complexity, which is of primary concern for real-time model predictive control.

Figure 2.20b depicts the tank water height validation results. The red dotted line is the measurement of the height of the thermocline, and the black solid line is the corresponding output of the tank model. The tank model successfully captures the dynamics of the thermocline in the tank.

2.6 Cooling Tower

2.6.1 Model

The following steady-state cooling tower model employs a CoolTools/York regression-based model of the cooling tower approach temperature $T_{app} = T_{cws} - T_{wb}$. The cooling tower uses variable speed fans to track a leaving water temperature set-point and reject the effects of

changing wetbulb temperature. The polynomial formula is a function of the water flow rate, fan power, wet bulb temperature and cooling range.

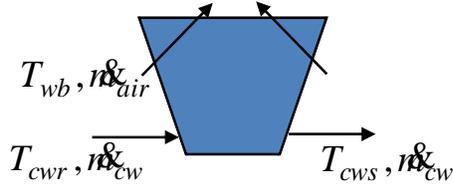


Figure 2.21: Cooling tower diagram

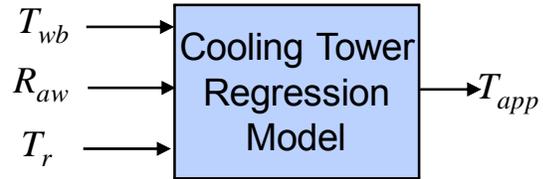


Figure 2.22: Cooling tower functional block diagram

The polynomial regression for the approach temperature

$$\begin{aligned}
 T_{app} = f_{CT}(R_{aw}, T_{wb}, T_R) @ & a_0 + a_1 T_{wb} + a_2 T_{wb}^2 + a_3 T_r + a_4 T_{wb} T_r + a_6 T_{wb}^2 T_r + a_7 T_r^2 + a_8 T_{wb} T_r^2 \\
 & + a_9 T_{wb}^2 T_r^2 + a_{10} R_{aw} + a_{11} T_{wb} R_{aw} + a_{12} T_{wb}^2 R_{aw} + a_{13} T_r R_{aw} + a_{14} T_{wb} T_r R_{aw} + a_{15} T_{wb}^2 T_r R_{aw} \\
 & + a_{16} T_r^2 R_{aw} + a_{17} T_{wb} T_r^2 R_{aw} + a_{18} T_{wb}^2 T_r^2 R_{aw} + a_{19} R_{aw}^2 + a_{20} T_{wb} R_{aw}^2 + a_{22} T_r R_{aw}^2 \\
 & + a_{23} T_{wb} T_r R_{aw}^2 + a_{24} T_{wb}^2 T_r R_{aw}^2 + a_{25} T_r^2 R_{aw}^2 + a_{26} T_{wb} T_r^2 R_{aw}^2 + a_{27} T_{wb}^2 T_r^2 R_{aw}^2
 \end{aligned} \quad (2.20)$$

is a function of the air to water ratio $R_{aw} = (\dot{m}_{cw} / \dot{m}_{cw,0}) / (N_{fan} / N_{fan,0})$, the range temperature $T_R = T_{cwr} - T_{cws}$ and the wet bulb temperature T_{wb} . The power consumed by the fan drive is represented by a cubic function $P_{fan} = P_{fan,0} (N_{fan} / N_{fan,0})^3$.

notation	description
N_{fan}	fan speed (rpm)
$N_{fan,0}$	design fan speed (rpm)
P_{fan}	fan power (kW)
$\dot{m}_{cw,0}$	design water mass flow rate
R_{aw}	air to water ratio
T_{app}	approach temperature [C]
T_R	range temperature [C]
T_{wb}	wet bulb temperature [C]
\cdot_{cwr}	cooling tower inlet
\cdot_{cws}	cooling tower outlet

2.6.2 Model Validation

The cooling tower at the University of California Merced Central Plant consists of five identical cells, with each cell possessing its own variable speed fan to regulate the exiting condenser water supply temperature. For the purpose of calibration, all five cell towers were treated as a single cooling tower by averaging the fan speeds of the five cell towers.

Initially, the parameters for the cooling tower calculation were calibrated using nearly two years of historical data from the University of California Merced. Since the condenser loop has constant speed pumps the water mass fraction is constant, $\dot{m}_{cw} / \dot{m}_{cw,0} = 1$, and a least squares fit was used to determine parameters a_i . However, it was observed that approximately the same quality of fit could be achieved by adjusting the ratio $\dot{m}_{cw} / \dot{m}_{cw,0}$ and employing the parameters a_i specified in [17]. Thus the former, more thoroughly documented, case was used for our purposes. As shown in Figure 2.23, the calibrated cooling tower model is capable of predicting the leaving cooling tower temperature within $\pm 1K$. The quality of the model prediction over a five nights of operation is shown in Figure 2.25.

Table 2.5: Calibrated cooling tower parameters

Name	Value	Name	Value	Name	Value
$\dot{m}_{cw} / \dot{m}_{cw,0}$	1.55	$N_{fan,0}$	100%	$P_{fan,0}$	95 kW
a_1	3.5974×10^{-1}	a_{10}	2.8476	a_{19}	-6.5376×10^{-2}
a_2	-5.5053×10^{-2}	a_{11}	-1.2176×10^{-4}	a_{20}	-2.2781×10^{-3}
a_3	2.3850×10^{-3}	a_{12}	1.4599×10^{-3}	a_{21}	2.5002×10^{-4}
a_4	1.7392×10^{-1}	a_{13}	1.6804	a_{22}	-9.1056×10^{-2}
a_5	-2.4847×10^{-2}	a_{14}	-1.6692×10^{-2}	a_{23}	3.1817×10^{-3}
a_6	4.8430×10^4	a_{15}	-7.1905×10^{-4}	a_{24}	3.8621×10^{-5}
a_7	-5.58984×10^3	a_{16}	-2.5485×10^{-2}	a_{25}	-3.4285×10^3
a_8	5.7700×10^4	a_{17}	4.8749×10^{-5}	a_{26}	8.5658×10^{-6}
a_9	-1.3424×10^{-5}	a_{18}	2.7192×10^{-6}	a_{27}	-1.5168×10^{-6}

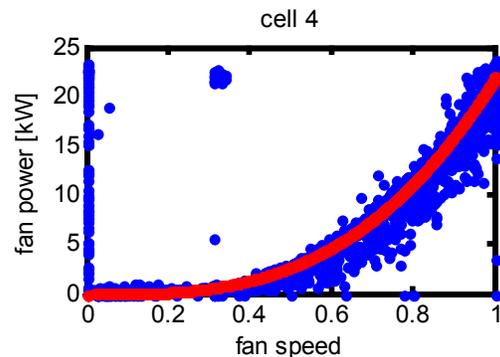
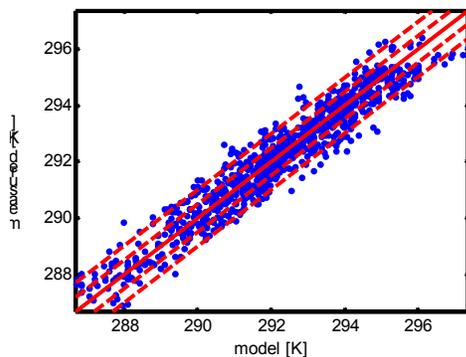


Figure 2.23: Measured versus model predicted leaving water temperature

Figure 2.24: Validation of cubic relationship between fan speed and power

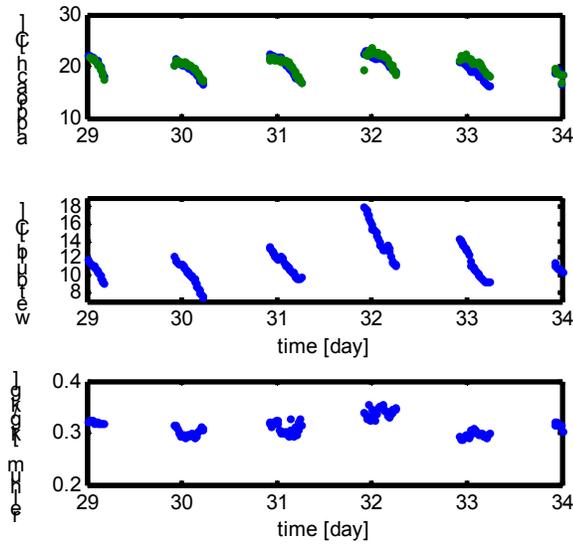


Figure 2.25: Five day time plot of measured versus predicted approach temperature

2.7 Campus Load Model

2.7.1 Model Description

The objective of the campus load model is to predict the campus cooling load based on the ambient conditions, time of day, day of the week, and time of year. The building load model has two main sub-components: a Gains component and a BuildingThermalLoad component.

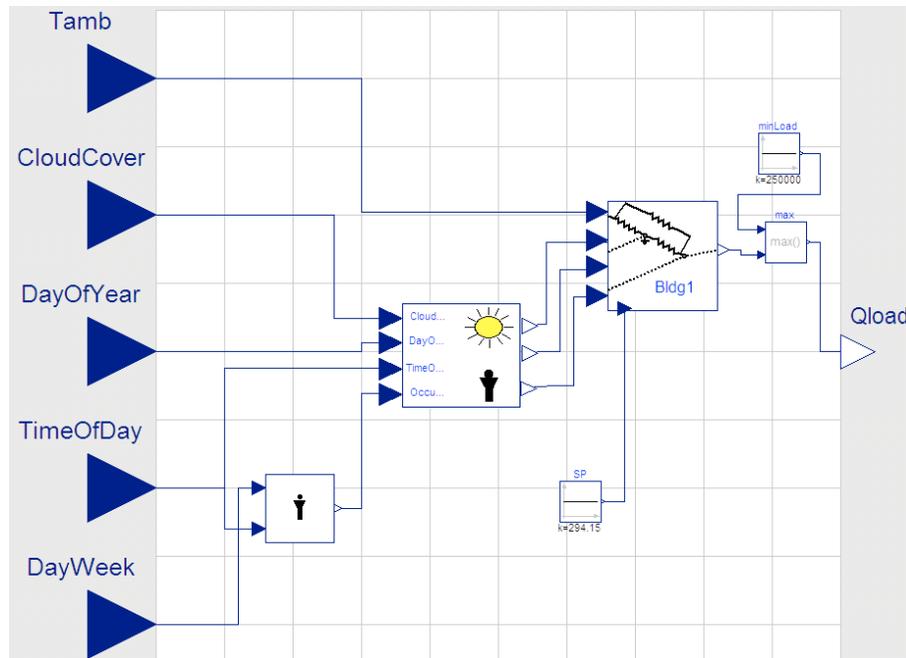


Figure 2.26: Building load model

notation	description
occDayStart:	hour of the day when the daytime internal load starts
occDayEnd:	hour of the day when the daytime internal load ends
QintBaseline:	a constant load independent of occupancy, ie. the overnight load
QintOccupied:	increase in internal load during occupied times on weekdays
QintOccupiedSaturday:	increase in internal load during occupied times on Saturdays
QintOccupiedSunday:	increase in internal load during occupied times on Sundays
bldgSolarOutsideQperExtratHoriz:	a building-geometry parameter, solar gain on the outer wall mass per unit of clear-sky horizontal radiation
bldgSolarInsideQperExtratHoriz:	a building-geometry parameter, solar gain on the inner wall mass per unit of clear-sky horizontal radiation
cloudCover:	impact of cloud cover (1=clear sky, 0=solar radiation completely blocked)

GainsConverter component

The GainsConverter component takes the time of day, day of week, day of year and cloud cover as its inputs and outputs a solar load and an internal load. The model currently has 12 parameters whose values are determined by calibration (listed as follows), and one parameter (site latitude = 37.302° for Merced) that is specified by the user. The internal load is essentially just a scheduled heat addition to the space, representing the heat gains from people, lights and equipment, with different values for nighttime, daytime during the week and daytime on Saturday and daytime on Sunday. The internal gain is parameterized by occDayStart, occDayEnd, QintBaseine, QintOccupied, QintOccupiedSaturday, and Qint Occupied Saturday.

The solar load is determined by first calculating the extraterrestrial horizontal radiation (an atmospheric-independent value which is a function of the time of day and day of year), and then using the cloud cover and two parameters for the building's solar gain susceptibility to calculate the inside and outside solar gain values. The solar gain is parameterized by bldgSolarOutsideQperExtratHoriz, bldgSolarInsideQperExtratHoriz, and cloudCover.

The inside and outside solar gain are calculated as follows

$$\begin{aligned} Q_{\text{solarOutside}} &= \text{cloudImpact} \cdot \text{extraterrHorizSolar} \cdot \text{bldgSolarOutsideQperExtratHoriz} \\ Q_{\text{solarInside}} &= \text{cloudImpact} \cdot \text{extraterrHorizSolar} \cdot \text{bldgSolarInsideQperExtratHoriz} \end{aligned}$$

where

$$\begin{aligned} \text{solarHour} &= (\text{TimeOfDay} - 12) \cdot 15^\circ \\ \text{solarDec} &= -23.45^\circ \cdot \cos\left(360^\circ \cdot \frac{(\text{DayOfYear} + 10)}{365}\right) \end{aligned}$$

$$\text{extraterrHorizSolar} = \max\left(0, \cos(\text{solarHour}) \cos(\text{solarDec}) \cos(\text{latitude}) + \sin(\text{solarDec}) \sin(\text{latitude})\right)$$

BuildingThermalLoad component

The BuildingThermalLoad component is based on the thermal resistance and capacitance model shown in Figure 12. (Various lumped parameter models like this have been developed and used in the literature. See Lee and Braun (2004) for a good example and further references.) As inputs, it takes the ambient temperature (θ_{outside}), the outside solar load (q_{solarOut}), the inside solar load (q_{solarIn}), the internal load (q_{internal}) and the indoor temperature set-point. It outputs the cooling load (q_{load}), and stores the temperature state of the thermal mass (T_{mass}). The component has six parameters: UAeff_inst ($= \frac{1}{R_1}$), UAc_outside ($= \frac{1}{R_2}$), UAc_inside ($= \frac{1}{R_4}$), UAc_middle ($= \frac{1}{R_3}$), WallCapInside ($= C_{\text{massIn}}$), and WallCapOutside ($= C_{\text{massOut}}$).

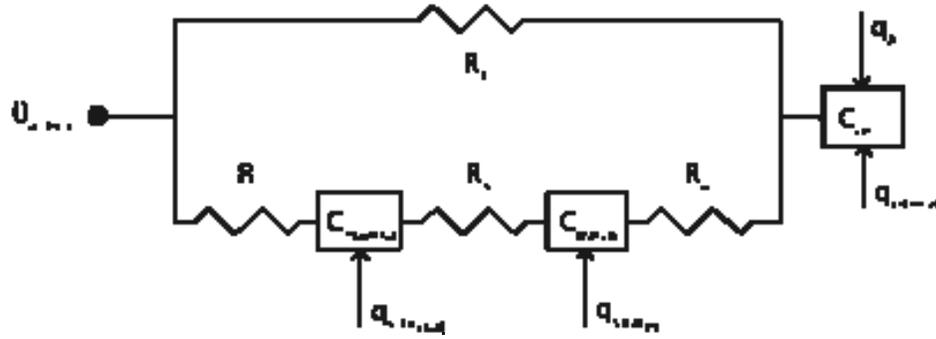


Figure 2.27: Thermal resistance and capacitance model

$$q_{load} = \max \left(0, q_{internal} + \frac{1}{R_4} (T_{massIn} - T_{zone}) + \frac{1}{R_1} (T_{amb} - T_{zone}) \right) \quad (2.21)$$

$$\dot{T}_{massIn} = \frac{q_{solarIn} + \frac{1}{R_3} (T_{massOut} - T_{massIn}) + \frac{1}{R_4} (T_{zone} - T_{massIn})}{C_{massIn}} \quad (2.22)$$

$$\dot{T}_{massOut} = \frac{q_{solarOut} + \frac{1}{R_2} (T_{amb} - T_{massOut}) + \frac{1}{R_3} (T_{massIn} - T_{massOut})}{C_{massOut}} \quad (2.23)$$

Overall inputs, outputs and parameters

The model has the following inputs, outputs and parameters:

Inputs

- T_{amb} , ambient temperature, in $^{\circ}\text{C}$
- $CloudCover$, a value between 0 (clear) and fully overcast (1)
- $DayOfYear$
- $DayOfWeek$, with Monday = 1
- $TimeOfDay$, hour of day
- T_{SP} , temperature set-point for the building, in $^{\circ}\text{C}$ (held constant)

Outputs

- Q_{load} , the cooling load for the building, in W

Parameters

- Internal Gain
 - occDayStart, occDayEnd
 - QintBaseline, QintOccupied, QintOccupiedSaturday, QintOccupiedSunday
- Building Thermal Properties
 - UAeff_inst, UAc_outside, UAc_inside, UAc_middle
 - WallCapOutside, WallCapInside

2.7.2 Model Calibration

The model was calibrated with annual data during pre-experiment development, then re-calibrated right before each experiment with the two previous weeks to ensure that it was up to date with any changes in occupancy or other attributes. The following shows the calibrated parameter values and performance for the first experimental period.

Table 2.5: Campus load model calibrated parameter values

Name	Value	Name	Value
cloudCover.fewClouds	0.175	QintOccupiedSaturday	250000
cloudCover.scatteredClouds	0.4	QintOccupiedSunday	125000
cloudCover.brokenClouds	0.6	QintBaseline	527000
cloudCover.overcast	0.8	UA_inst	66200
bldgSolarOutsideQperExtratHoriz	400000	UAcOut	100000
bldgSolarInsideQperExtratHoriz	400000	UAcIn	100000
occDayStart	7	UAcMiddle	99600
occDayEnd	20	WallCapInside	5000012800
QintOccupied	500000	WallCapOutside	500000000

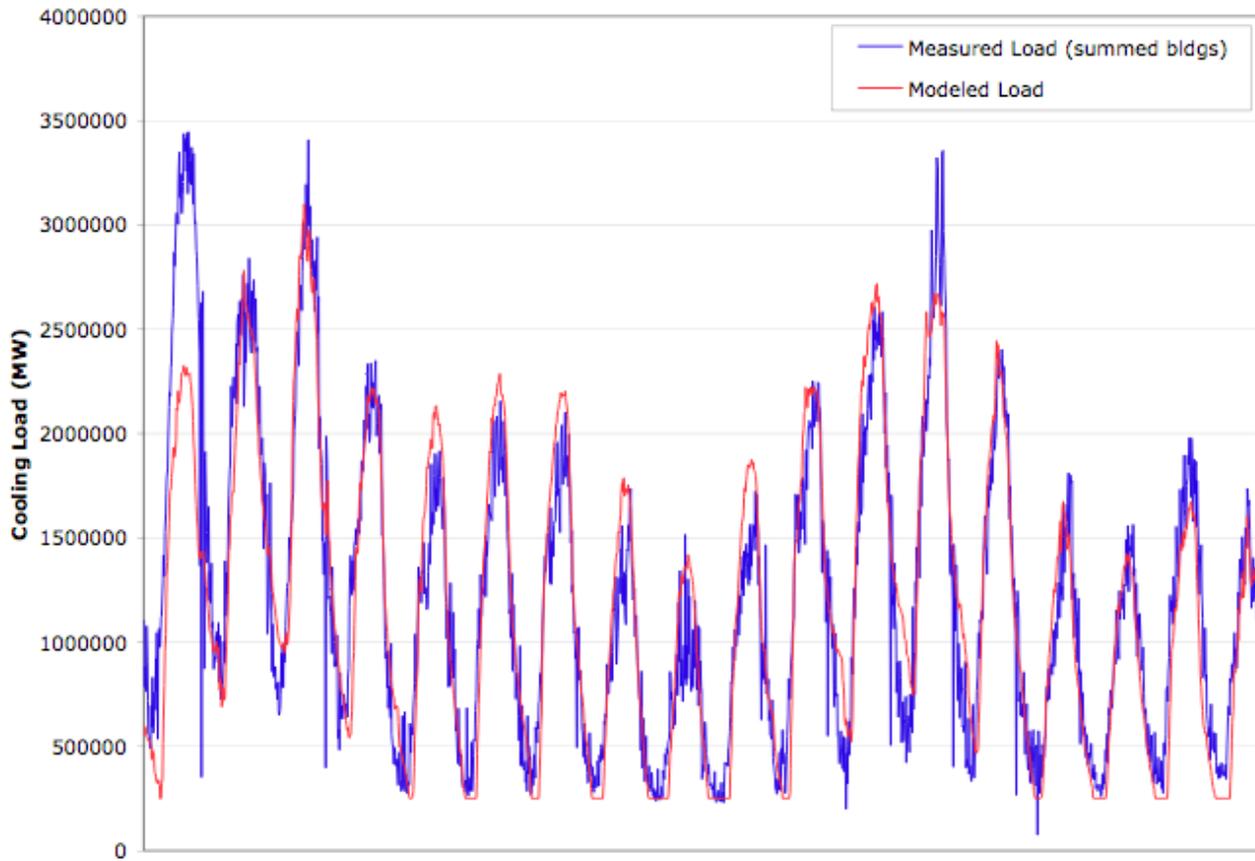


Figure 2.28: Campus cooling load, May16–Jun2, 2009

2.8 Modeling the Campus Chilled Water Return Temperature

2.8.1 Model Description

The interaction of the campus buildings with the chilled water system is divided into two parts. The load is modeled as a function of weather and occupancy disturbances, as discussed in the next section. It is not dependent on any other models for its input values, and thus may be treated as a separate disturbance model whose values can be pre-calculated at each controller time step and does not need to be repeated iteratively within the online optimization. The flow rate and return temperature from the campus is a function of this load and the ambient temperature, but is also a function of the chilled water supply temperature, which is a control variable, so the return temperature must be modeled separately and must be included in the online system model.

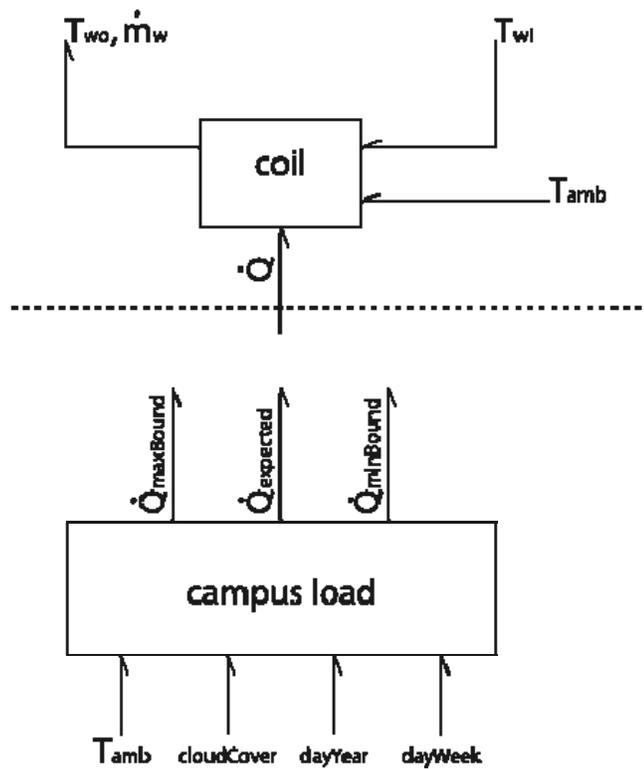


Figure 2.269: Model Context

notation	description
\dot{m}_w	water mass flow rate [kg/s]
\mathcal{Q}_{load}	cooling load from building model [W]
$T_{airReturn}$	average return air temperature [K]
T_{ai}	mixed air temperature [K]
T_{amb}	ambient air temperature [K]
T_{ao}	average supply air temperature [K]
T_{wi}	water supply temperature [K]
T_{wo}	return water temperature [K]
UA	heat transfer effectiveness [W/K]
α_1, α_2	curve fit parameters for steady state UA
β	curve fit parameter for UA decrease at low loads due to control instability
γ	outside air fraction
$\xi(\cdot)$	building-specific function for control instability as a function of load (\mathcal{Q}_{load})

Table 2.6: Campus equivalent cooling coil model calibrated parameter values

Constants	$c_{p_a} = 1012 \text{ W/kg.K}$, $c_{p_w} = 4181.3 \text{ W/kg.K}$
Equations	<i>Explicit:</i>
	$T_{ai} = T_{airReturn} + \gamma(T_{amb} - T_{airReturn})$
	$\dot{m}_a = \dot{Q}_{load} / (c_{p_a} (T_{ai} - T_{ao}))$
	$\zeta = (1 + \beta \zeta^2)^{-1}$
	$\dot{Q}_{load} = \dot{m}_w c_{p_w} (T_{wo} - T_{wi})$
	$\dot{Q}_{load} = UA\Delta T_m = \zeta (\alpha_1 \dot{m}_w^{0.85} + \alpha_2 \dot{m}_a^{0.95})^{-1} \frac{T_{ao} - T_{ai} - T_{wi} + T_{wo}}{\log\left(\frac{T_{ao} - T_{wi}}{T_{ai} - T_{wo}}\right)}$

Model Principles

The model is used to predict the return water temperature from the campus. It must be calibrated to fit the measured data for the campus, but must also have a reasoned basis in order to predict the effects of operating conditions outside of the historical data, particularly in terms of changes in the supply water temperature. The historical data, which is predominantly with supply water temperatures close to 4°C (39°F), shows the following trends: the return temperature tends to reduce and become more scattered as the load decreases.

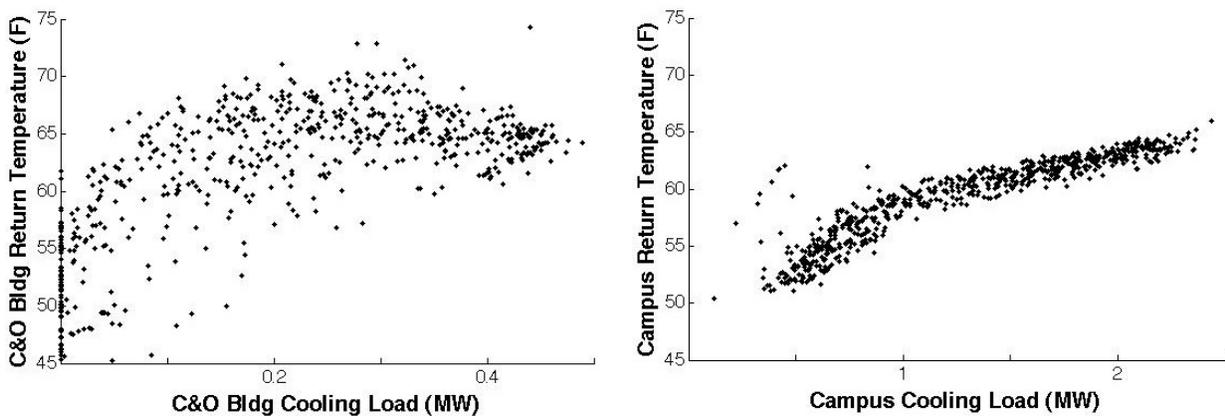


Figure 2.30: Measured Return Temp, Aug23-Sep4 2009 for the Classroom and Office Building and for the entire Campus

The first-principle equations used as the basis for this model are

$$\dot{Q}_{load} = \dot{m}_w c_{p_w} (T_{wo} - T_{wi}) \quad (2.25)$$

$$\dot{Q}_{load} = UA\Delta T_m \quad (2.26)$$

where the Log Mean Temperature Difference, ΔT_m , is defined as follows

$$\Delta T_m = \left(\frac{T_{ao} - T_{ai} - T_{wi} + T_{wo}}{\log \left(\frac{T_{ao} - T_{wi}}{T_{ai} - T_{wo}} \right)} \right) \quad (2.247)$$

Variation of UA as a function of flow rates and load

A constant UA does not properly explain the behaviour shown in Figure 2.30. Based on the work of Holmes [19], the following relationship of the overall UA to the air and water flow rates is used, where α_1 and α_2 are calibration parameters:

:

$$UA_{theoretical} = (\alpha_1 \dot{m}_w^{-0.85} + \alpha_2 \dot{m}_a^{-0.85})^{-1} \quad (2.28)$$

This provides a physical basis for the model, but it does not adequately capture the reduction in UA at low loads, and thus cannot be calibrated to fit the data very well. As shown in Figure 2.31, a model based only on these equations, regardless of its calibration, tends to produce a relatively constant return temperature.

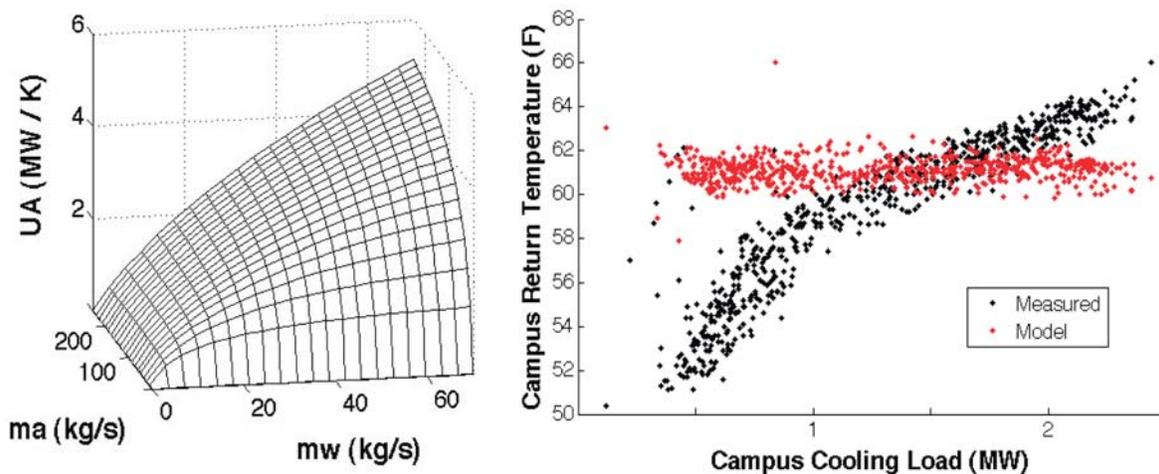


Figure 2.31: UA vs water and air flow rates (left), and Return Temp vs Load (right)

An Empirical Addition to the Model to Account for the Reduction in UA at Low Loads

There are various possible reasons why the actual UA is lower at low loads than predicted by these equations. One such reason is discussed in the next subsection: greater flow rate fluctuations around set-points are found in the building heat exchangers at lower loads, and since the UA vs flow rate curve is concave, this could cause the effective (average across fluctuations) UA to be lower than the theoretical, as well as possibly causing various transient effects that may also decrease the effective UA. Further analysis of this and other possible causes is currently underway as a follow-up to this research, but for the purpose of this MPC case study, a simplified semi-empirical model with fast run-times is what was required, so an empirical calibration factor was added to the model to account for the observed UA reduction at low loads. This addition, shown below, was designed to retain the first-principles basis of the model as much as possible, while still allowing for a better fit to the empirical data.

$$UA = \zeta (UA_{theoretical}) \quad (2.29)$$

$$\zeta = (1 + \beta \xi^n)^{-1} \quad (2.30)$$

The variable ξ is an empirically-defined logarithmic curve. For the MPC experiment, the curve used and its rationale is described in the following subsection. The variable β is a calibration parameter between 0 and 1 that determines the magnitude of the deviation from the theoretical. With this adjustment, the UA at low loads can be decreased, as shown in the left part of

Figure 2. 3232, and the model can be calibrated to fit the measured data much more accurately (as shown in the right part of

Figure 2. 3232).

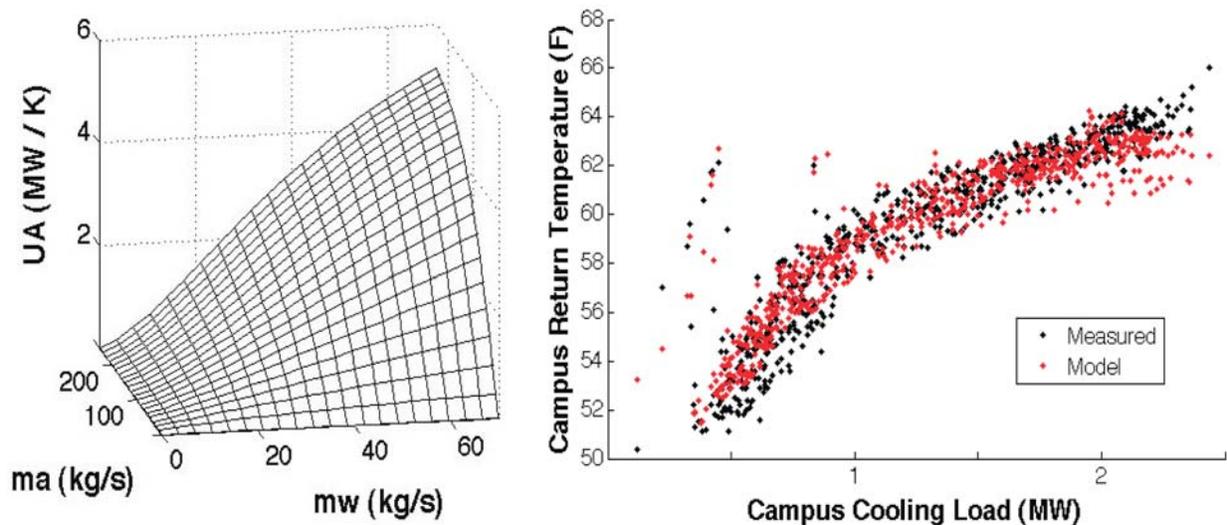


Figure 2. 32: UA vs water and air flow rates (left), and Return Temp vs Load (right), after empirical UA reduction factor added to model

Temperature Oscillations at Building Heat Exchangers at Low Loads [bmh7]

Error! Reference source not found.2.33 shows the difference between the supply air temperature and the set-point versus load for the Classroom and Office Building. The 15-minute sampled data in **Error! Reference source not found.**2.33 exhibit substantial fluctuations of the supply air temperature about its set-point at low load, whereas the two-hour averaged data (bottom graph) show much smaller fluctuations, suggesting a control stability problem at low loads. The outliers with large negative errors were from a two-week problem in May and June, 2009, provisionally attributed to a valve fault.

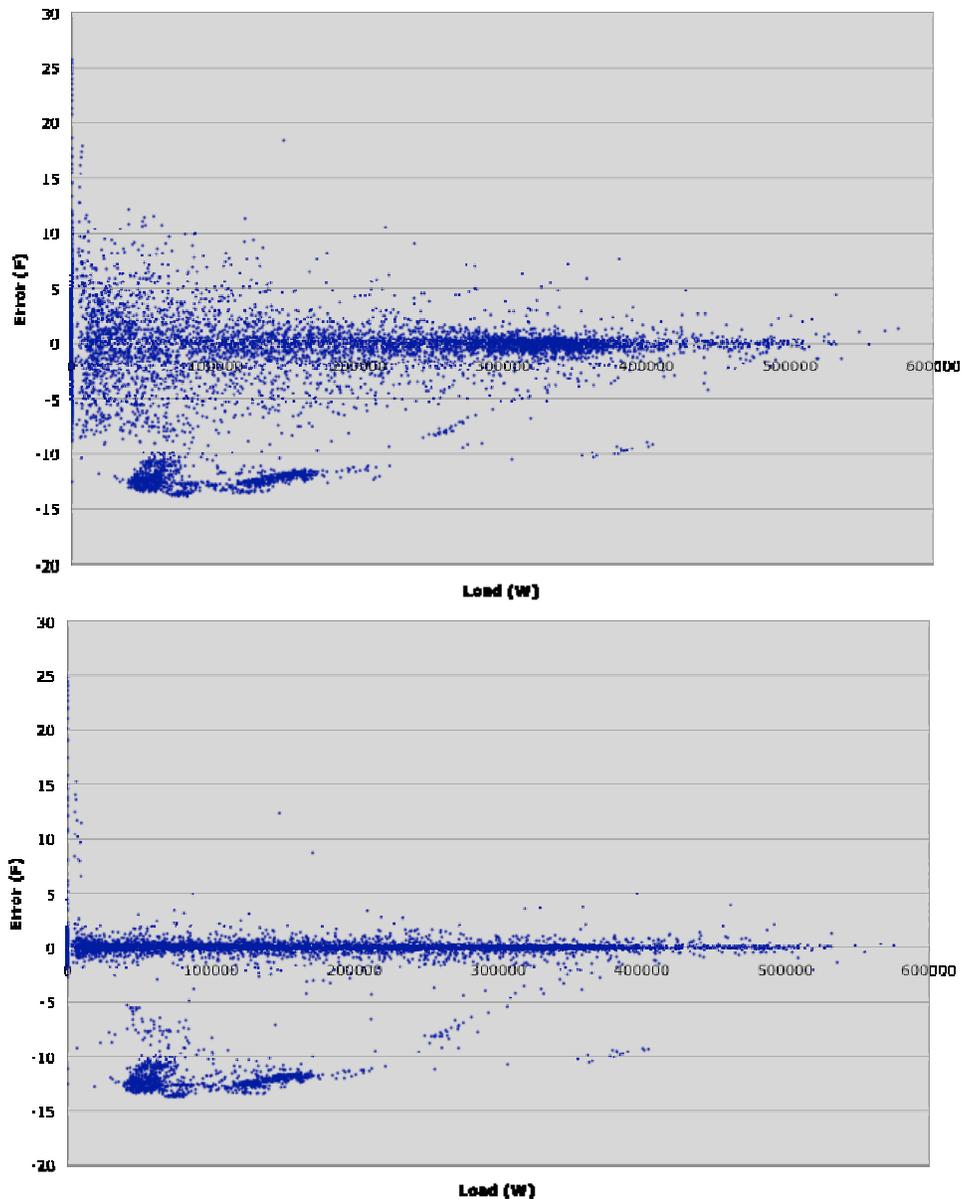


Figure 2.33: Supply Air Temperature control error, 15-minute samples (top), 2-hr averages (bottom)

The supply air temperature oscillations at low cooling loads are likely caused by the feedback control. The supply air temperature is regulated by modulating the chilled water flow rate through the cooling coil via a control valve. A proportional integral (PI) controller adjusts the control valve based on the measured error between the supply air temperature and its set point. There are two key differences in the cooling coil and valve behavior at low loads versus higher loads. First, the valve is nearly closed. As a result of non-linearities in the valve characteristic, small increments in the valve opening cause larger changes in flow rate than when the valve is more open. Second, a small variation in the T_{wo} has a larger impact on the log

mean temperature difference as T_{ai} approaches T_{wo} . Those two differences effectively change the gain of the system, causing oscillations in the feedback loop tuned for higher cooling loads.

As noted above, this oscillation at low loads is one of the possible causes of the lower than expected UA at low loads. The magnitude of these oscillations was used in creating the empirical curve for ξ . The following equation provides a relationship between the oscillations (absolute value of the error) and the load, as shown in

Figure 2.27, and is of a form that can readily be used in the empirical UA modification equation above. (Adding 1 inside the logarithm keeps it well-defined for all positive loads, and normalizing it between 0 and 1 within the operating range makes it easier to deal with.)

$$\xi = \frac{-0.9074 \ln(1 + \frac{Q_{load}}{12.329}) + 12.329 - 0.3}{12.329 - 0.3} \quad (2.31)$$

Equation 2.31 was used in the model for the MPC. Its relationship with a potential physical cause of the UA reduction was seen as beneficial, but the main reason for its selection is the ability of the model to capture the observed return temperature behavior when this equation is used. (Note that, in hindsight, the authors acknowledge that even though useful, this approach is somewhat intellectually unsatisfying, and that future iterations of the model should either have a more direct link between physical UA-reduction causes and the model equations, or it should use a more transparent empirical approach to ζ (and ξ , if still applicable).)

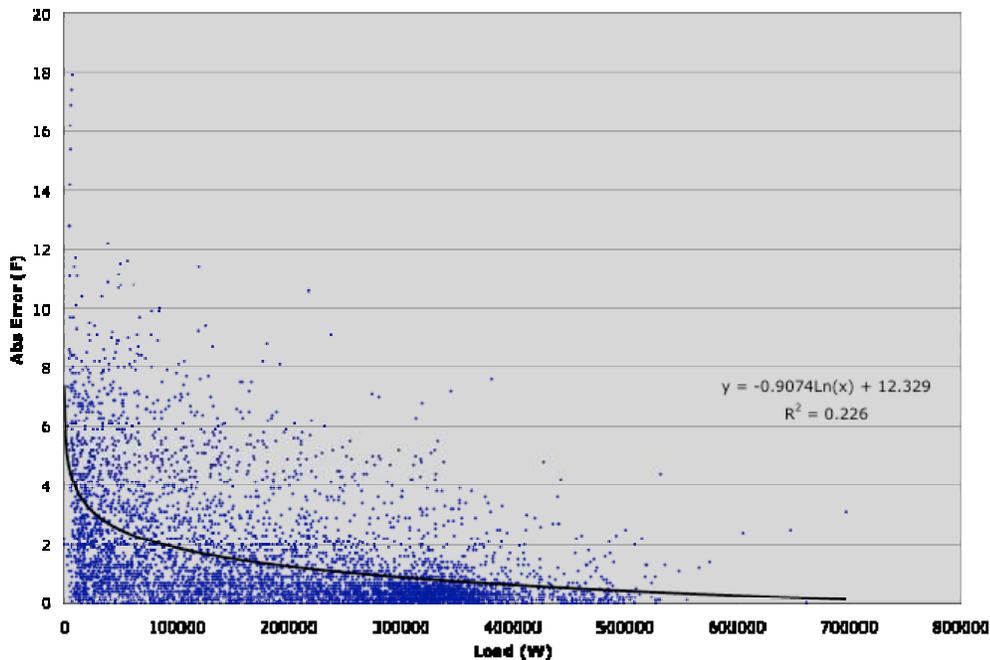


Figure 2.27: Absolute Error vs Load

2.8.2 Model Calibration

The four model parameters, α_1 , α_2 , β and γ , were calibrated as follows, using approximately one month of 15-minute data:

1. Determine $T_{\text{airReturn}}$ and T_{ao} as averages of measured data. Fit γ using the determined value of $T_{\text{airReturn}}$ and the data for T_{amb} and the mixed air temp T_{MA} .
2. Fit the parameters α_1 and α_2 with $\beta=0$ and only using points where the load is greater than some heuristically-determined value (1.5MW for the campus)
3. Fit the parameter β using the α_1 and α_2 values from step 2 and all of the data

Table 2.7 shows the calibrated values for the model, and

Figure 2. 32b (above) compares the model to measured data under normal operation. Appendix 3 describes a slight variation that was made to the model after a campus supply water temperature experiment. This experiment and variation did significantly not affect the model or the controller, but may be of interest in future model development and thus is described in detail in the appendix.

Table 2.7: Calibrated Parameter Values for the Merced Campus

$T_{\text{airReturn}} = 296.66 \text{ K}$	$\gamma = 0.13551$
$T_{\text{ao}} = 284.84 \text{ K}$	$\beta = 56.025$
$\alpha_1 = 7.3443\text{e-}05$	$\alpha_2 = 1.5464\text{e-}04$

2.9 Interconnection of Detailed Models

The following discusses the interconnection of the detailed component models discussed preceding this subsection. Specifically, the following details the implicit mass and energy balances for the chilled water plant illustrated in Figure 2.1. As aforementioned, the chilled water plant rotates among the three chillers, and the evaporators of the two operational chillers are piped in series. Consequently, the two operational chillers will be differentiated by referring to each chiller as either the upstream or downstream chiller. Beginning with the condenser loop, as shown in Figure 2.1, the condenser water mass balance is

$$\dot{m}_{\text{cw}} = \dot{m}_{\text{cw,ch1}} + \dot{m}_{\text{cw,ch2}}, \quad (2.25)$$

the return condenser water energy balance is

$$\dot{m}_{\text{cw}} h_{\text{cwr}} = \dot{m}_{\text{cw,ch1}} h_{\text{cwr,ch1}} + \dot{m}_{\text{cw,ch2}} h_{\text{cwr,ch2}}, \quad (2.26)$$

and the supply condenser water specific enthalpy is

$$h_{cws} = h_{cws,ch1} = h_{cwr,ch2}, \quad (2.27)$$

where $\dot{m}_{cw,ch1}$ and $\dot{m}_{cw,ch2}$ are flow rates through the upstream and downstream chillers, respectively. Two fixed speed pumps supply condenser water to a common header, and the common header supplies condenser water to the two chiller condenser also piped in parallel. As confirmed by individual condenser flow-rate measurements, the flow-rates through each condenser are approximately equal, i.e. $\dot{m}_{cw,ch1} \approx \dot{m}_{cw,ch2}$. As observed from EMCS data, the flow rate through each condenser was approximately 3800 gpm (240 kg/s). The pumps run at fixed speed and the pressures losses should be constant due to the static configuration of the condenser loop. Based on the pump specification sheets, each condenser pump is modeled as consuming 60 kW when the chillers are operational.

The primary chilled water loop is formed as the circuit between the chiller and the chilled water storage tank. As shown in Figure 2.1, the chilled water plant operates with the evaporators of two chillers piped in series, thereby yielding the specific enthalpy and mass flow-rate relationships

$$h_{chws,ch1} = h_{chwr,ch2}, \quad \dot{m}_{chw} = \dot{m}_{chw,ch1} = \dot{m}_{chw,ch2}. \quad (2.28a, 2.30b)$$

As shown in Figure 2.1, the chilled water supplied to the chillers is a blend of water from the thermal energy storage tank and the secondary (campus) chilled water loop, thus the mass balance is

$$\dot{m}_{chw} = \dot{m}_{tank} + \dot{m}_{campus} \quad (2.29)$$

and the energy balance is

$$\dot{m}_{chw} h_{chwr,ch1} = \dot{m}_{tank} h_a + \dot{m}_{campus} h_{chwr,campus}. \quad (2.30)$$

Similarly, the chilled water supplied to the campus is a blend of water from the chillers and the thermal energy storage tank, thus the mass balance is

$$\dot{m}_{campus} = \dot{m}_{chw} - \dot{m}_{tank} \quad (2.31)$$

and the energy balance is

$$\dot{m}_{campus} h_{chws,campus} = \dot{m}_{chw} h_{chws,ch2} - \dot{m}_{tank} h_b. \quad (2.32)$$

Note the reason why (2.29) and (2.32) represent the same relationship is because chilled water plant consists of closed chilled water loops. For example, the combined mass of water in

the thermal energy storage tank is constant and the flow supplied to the secondary (campus) chilled water loop is the same as the flow-rate returning from the secondary (campus) chilled water loop.

2.10 Baseline Policy Model

The baseline control policy constitutes the control logic and operation procedures typically followed by the central plant staff while operating the chilled water plant. Two baseline policies will be detailed in the following discussion. Baseline policy #1 describes the policy as observed prior to the first experimental implementation of the MPC algorithm in June 2009, whereas baseline policy #2 describes the policy after the first MPC experiment.

As represented by the numerous K_{PI} blocks in Figure 2.1, a number of proportional-integral (PI) controllers are implemented in the energy management control system software. The function of these blocks is to ensure tracking of set-points specified in the EMCS software. Since the dynamic responses of the control loops are significantly faster than the dynamics of interest in the study, the control loops are modeled as ideal, i.e. perfect set-point following subject to capacity limitations.

The chilled water temperature supply set-point for chiller 2 was set at 3.9°C (39°F), the temperature at which water has its maximum density, in order to maximize buoyancy driven stratification in the thermal storage capacity of the tank. In order to distribute the cooling load equally between the two chillers, the chilled water supply temperature set-point for chiller 1 $T_{chws,ch1}$ is the average of the temperature of the water entering the first chiller, $T_{chwr,ch1}$, and the temperature of the water leaving the second chiller, $T_{chws,ch2}$.

The primary loop chilled water flow-rate set-point $\dot{m}_{ch,ref}$ is scheduled based on the chilled water return temperature entering the first chiller $T_{chwr,ch1}$. When $T_{chwr,ch1}$ is above 18.3°C (65°F), the pump speed is modulated to maintain the minimum flow-rate at 145 kg/s (2310 gpm). Conversely, when $T_{chwr,ch1}$ is less than 12.8°C (55°F), the pump speed is modulated to attain the maximum design flow-rate of 220 kg/s (3500 gpm). The objective of scheduling the flow-rate based on the chilled water return temperature is two-fold. First, the flow-rate through the chillers is reduced in response to increases in $T_{chwr,ch1}$ in order to avoid exceeding the cooling capacity of the chillers. Conversely, the flow-rate through the chillers is increased in response to a decrease in $T_{chwr,ch1}$ in order to maintain loading on the chillers. As shown using the calibrated chiller curves, see Figure 2.9, maintaining ideal part loading on the chillers is critical for chiller efficiency. The maximum physically attainable flow-rate was 183 kg/s (2900 gpm) before June 2009. After advising the plant operators that this limitation in the flow-rate had significantly reduced the chilled water plant COP, the pump balancing valves were adjusted to obtain a maximum of 201 kg/s (3200 gpm) after the first MPC experiment.

The cooling tower set-point $T_{cws,ref}$ is scheduled to increase with the current drawn by the chillers, expressed as a percentage of the full load amps (%FLA). As discussed in the cooling

tower model, the fan speed of the cooling tower is modulated to track the set-point $T_{cws,ref}$. According to the control logic for the EMCS, $T_{cws,ref}=58^{\circ}\text{F}$ when %FLA is less than 50% and $T_{cws,ref}=60^{\circ}\text{F}$ when average %FLA is greater than 90% for baseline policy #1. In contrast, the original sequence of operations specified $T_{cws,ref}=68^{\circ}\text{F}$ when the average %FLA is greater than 90%. In the first MPC experiment, it was observed that it was advantageous to increase $T_{cws,ref}$ beyond the $_{[PH9]}60^{\circ}\text{F}$ limit prescribed by baseline #1. Subsequently, the operators increased the upper-limit to 68°F in baseline #2.

2.10.1 Weather Forecasts

In order to obtain weather forecasts for use by the model predictive control algorithm, the National Digital Forecast Database (NDFD) was utilized. A Simple Object Access Protocol request was used to query the NDFD and parse and pass the weather forecast for the University of California Merced to the model predictive control algorithm. At this stage, only the forecasted dry bulb temperature and relative humidity were used. In the future additional forecasted variables, such as sky cover, could be correlated with the campus cooling load prediction.

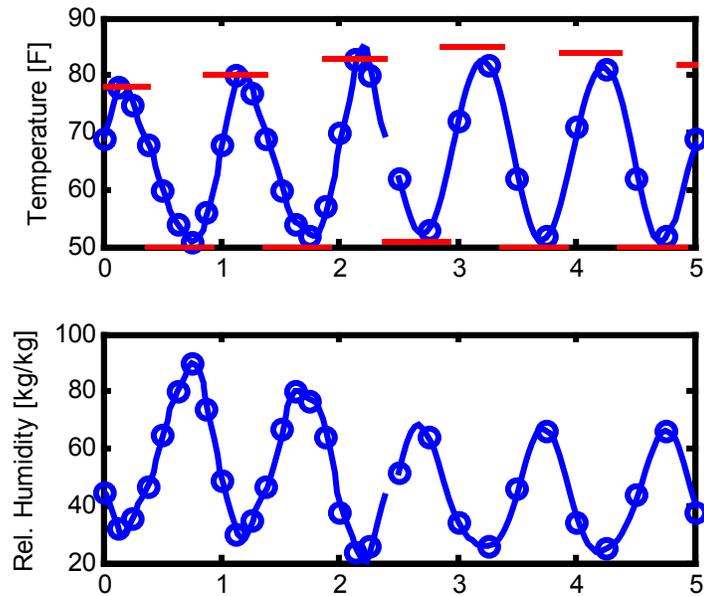


Figure 2.28: Sample dry bulb temperature and relative humidity forecast

2.10.2 Energy Pricing

UC Merced is currently enrolled in a special plan, electric schedule E-20. The customer's monthly charge for the service under Schedule E-20 is the sum of a customer charge and energy charges, and all the unit price varies depending on the period of time.

Table 2.8: Definition of time periods

SUMMER	Period A	(May 1st though Oct. 31st)
Peak	12:00--18:00	except holidays
Partial-peak	8:30--12:00 AND 18:00--9:30	except holidays
Off-peak	21:30--8:30 ALL DAY	Mon. through Fri. Sat., Sun, and holidays
WINTER	Period B	(Nov. 1st though Apr. 30st)
Partial-peak	8:30--21:30	except holidays
Off-peak	21:30--8:30 ALL DAY	Mon. through Fri. Sat., Sun, and holidays

Table 2.9 shows the definition of the time periods. The unit electricity prices defined in Table 2.9 are denoted by the function $C(t)$. There are no demand (kW) charges.

Table 2.9: Total Energy Rates

Period	Rates (\$/kWh)
Peak Summer	\$ 0.13593
Part-Peak Summer	\$ 0.09204
Off-Peak Summer	\$ 0.07392
Part-Peak Winter	\$ 0.08155
Off-Peak Winter	\$ 0.07118

3 Model Predictive Control Development

This section presents the design of a MPC controller whose objective is to find the optimal control sequence that satisfies the required cooling load and minimizes electricity usage. For complex constrained multivariable control problems, MPC has become the accepted standard in the process industries [20] its success is largely due to its unique ability to handle, systematically, simply and effectively, constraints on control and states. The scheme of the proposed MPC is depicted in Figure 3.1.

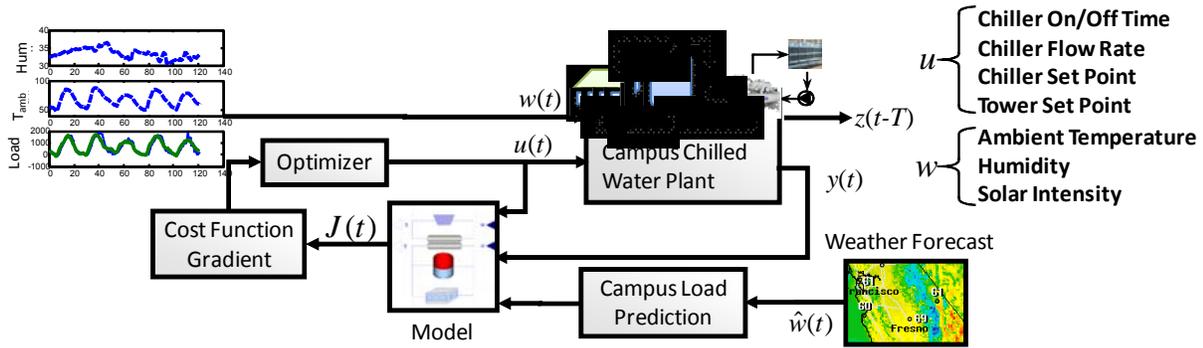


Figure 3.1 MPC Scheme

The idea of the proposed MPC scheme is to use the simplified models of the cooling plant and of the campus presented in the previous sections in order to predict the future evolution of the system. At each sampling time, starting at the current state and using weather predictions, an open-loop optimal control problem is solved over a finite horizon. The optimal command signal is applied to the process only during the following sampling interval. At the next time step a new optimal control problem based on new measurements of the state is solved over a shifted horizon. A typical open-loop finite time optimization problem is formulated as follows (details will be give in the next section)

$$\begin{aligned}
 J^\Sigma &= \min_{u_{t \rightarrow t+N|t}} P(x_{t+N|t}) + \sum_{k=0}^{N-1} L(x_{t+k|t}, u_{k+t|t}, k) & (3.1) \\
 \text{subj. to } & x_{t+k+1|t} = g(x_{t+k|t}, u_{k+t|t}, k, d_{t+k|t}), \forall k = 0, 1, \dots, N-1 \\
 & x_{t+k|t} \in \Xi, \forall k = 1, 2, \dots, N \\
 & u_{t+k|t} \in Y, \forall k = 0, 1, \dots, N-1
 \end{aligned}$$

where u are the controlled inputs (condensed water supply temperature, chilled water supply temperature and the chilled water supply mass flow rate) and x represent the system states (temperature and height of the cool water and warm water respectively) and d represent the

external disturbances (ambient temperature and campus load). $P(x)$ is the terminal cost and $L(x, u, t, d)$ is the stage cost. They are both used to capture the performance objective to be minimized (such as the electricity bill or the coefficient of performance (COP) of the central plant). The function $g(x, u, t, d)$ is a compact representation of the state update dynamic equation which describes the cooling plant and buildings. The control inputs u and the states x are subject to operational constraints which avoid system malfunctioning (such as chillers surging) and ensure that the campus cooling demand is satisfied over the prediction horizon. All variables use the following time indexing: $x_{t+k|t}$ denotes the state vector at time $t+k$ predicted at time t obtained by starting from the current state $x(t)$ and applying the input sequence, $u_{t \rightarrow t+N|t}$, and disturbance prediction, $d_{t \rightarrow t+N|t}$, (weather and campus load) to the system model.

The proposed MPC scheme minimizes a nominal performance index subject to a nominal model and constraints. It is not robust in the sense that neither input uncertainties nor model error are taken into account. This could be addressed by designing a robust MPC scheme [21]. A typical robust MPC scheme involves solving a min-max problem to optimize robust performance (the minimum over the control input of the maximum over the disturbance) while enforcing input and state constraints for all possible admissible bounded disturbances. Also, the proposed MPC in Equation 3.1 does not consider system faults. Robustness to faults can be indirectly obtained by detuning the controller (reducing the weights) at the price of a reduced performance or can be systematically and easily taken into accounts by switching to a different model $g(x, u, t, d)$ and constraints X and U when a fault is identified. The finite time optimization problem solved at UC Merced is detailed in the next section.

3.1 MPC details

The finite time optimization problem solved by the MPC is formulated by the following optimization problem:

$$J^{\Sigma}(x(t), t) = \min_{\hat{u}_{0|t}, \Lambda, \hat{u}_{M-1|t}, t_s, t_f} \sum_{i=1}^N \{ C(t+i\Delta t) E(x_{i\Delta t|t}, u_{(i-1)\Delta t|t}) \Pi_s + \Pi_{i\Delta t|t} \Pi_u \} \quad (3.2)$$

s.t.

$$y_{i\Delta t|t} \in \Psi, \forall i = 1, 2, \dots, \Lambda, N$$

$$u_{i\Delta t|t} \in Y, \forall i = 0, 1, \dots, \Lambda, N-1$$

$$y_{N\Delta t|t} \in \Psi_f(t)$$

$$E(x_{i\Delta t|t}, u_{(i-1)\Delta t|t}) = Power(x_{i\Delta t|t}, u_{(i-1)\Delta t|t}) \Delta t$$

$$[u_{0|t}, \Lambda, u_{(N-1)\Delta t|t}] = B \otimes I_m [\hat{u}_{0|t}, \Lambda, \hat{u}_{M-1|t}]$$

$$x_{(k+1)\Delta t|t} = f(x_{k\Delta t|t}, u_{k\Delta t|t}, \Phi_{k\Delta t|t}, k\Delta t), \forall k = 0, 1, \dots, \Lambda, N-1$$

$$y_{k\Delta t|t} = g(x_{k\Delta t|t}, u_{(k-1)\Delta t|t}, \Phi_{k\Delta t|t}, k\Delta t), \forall k = 1, 2, \dots, \Lambda, N$$

where $\Psi_f(t)$ is the terminal constraint set; $C(t)$ is the energy price defined in section 2.10.2;

$\Delta t = 1$ hour; S is the weight on the energy consumption; R_u is the weight of the control inputs; $Power(x, u)$ is the function to calculate the total power consumption of the central plant including the chillers and cooling towers.

Let $U_{1 \rightarrow N|t}^\Sigma = \{u_{1|t}^\Sigma, \Lambda, u_{N|t}^\Sigma\}$ be the optimal solution of Problem at time t , and $J_t^\Sigma(x(t))$ the corresponding value function. Then, the first element of $U_{1 \rightarrow N|t}^\Sigma$ is implemented to the system

$$u(t) = u_{1|t}^\Sigma \quad (3.3)$$

The optimization problem $J^\Sigma(x(t), t) = \min_{\hat{u}_{0|t}^\Sigma, \Lambda, \hat{u}_{M-1|t}^\Sigma, t_s, t_f} \sum_{i=1}^N \{ \mathbb{E}(t + i\Delta t) E(x_{i\Delta t|t}, u_{(i-1)\Delta t|t}) \Pi_S + \mathbb{E}_{i\Delta t|t} \Pi_{R_u} \}$

(3.2) is repeated at time $t+1$, with the updated new state $x_{0|t+1} = x(t+1)$, yielding a *moving or receding horizon control* strategy. The control sampling time is one day, and prediction horizon is set to 72h (one day).

The proposed MPC controller uses a move blocking strategy to reduce the computational time required for its real time implementation and a time varying terminal cost to ensure that the energy stored in the tank will be able to satisfy the capus cooling demand even in the worst case. Both strategies are detailed in the next sections.

3.2 Move Blocking Strategy

The prediction horizon of the proposed MPC controller is 24 hours, and the control sampling time is one hour. As a result, there are a total of 72 optimization variables, since the control input dimension is 3. It is common practice to reduce the degrees of freedom by fixing the input or its derivatives to be constant over several time steps [22]. Here, the Moving Window Blocking approach proposed in [23] is used. The following definitions are needed before providing the algorithm used.

Definition 1 (Admissible Blocking Matrix) A matrix $B \in \{0,1\}^{N \times M}$ is an admissible blocking matrix if $M < N$, and one entry in each row of B is equal to 1, the elements of the matrix are arranged in an "upper staircase" form, i.e. if the column in which a 1 occurs in the i 'th row is

$$j^\Sigma(i) := \{j \mid B_{i,j} = 1\}$$

then $j^\Sigma(i+1) \geq j^\Sigma(i)$ for all $i \in \{1, 2, \dots, N-1\}$. Where $B_{i,j}$ denotes the element of i 'th row and j 'th column of matrix B .

Definition 2 (Blocking Length Vector) Given an admissible blocking matrix $B \in \{0,1\}^{N \times M}$, the blocking length vector $L(B)$ is defined as the columnwise summation of the matrix B . An admissible block vector corresponds to a unique blocking length matrix.

The following algorithm is in the proposed MPC.

Algorithm 1 (Moving Window Blocking) Given an initial blocking length matrix L_0 ;

1. let $i = 0$;
2. let $\hat{L}_0 = [L_0, 0]$;
3. if $\hat{L}_i(1) > 1$, $\hat{L}_{i+1} := \hat{L}_i$; $\hat{L}_{i+1}(1) := \hat{L}_i(1) - 1$; $\hat{L}_{i+1}(end) := \hat{L}_i(end) + 1$.
if $\hat{L}_i(1) = 1$, $\hat{L}_{i+1} := [\hat{L}_i(2:end-1), \hat{L}_i(end) + 1, 0]$.
4. $L_{i+1} := L_{i+1}(1:end-1)$ if $L_{i+1}(end) = 0$ & L_{i+1} otherwise.
5. if $L_i = L_0$, stop. Otherwise, go to next step;
6. let $i := i + 1$, and go to step 3.

In Algorithm 1 $L_i(end)$ is the last element of L_i , and \hat{L}_i is the auxiliary vector. Algorithm 1 assumes that the system period is $T = N$ and generates a sequence of blocking length vectors $L_{k+N} = L_k$ which are obtained by properly shifting an initial choice L_0 at each time step in order guarantee persistent feasibility [6]. For instance, in this simple example, starting with $L_0 = [3, 1]$ yields $L_1 = [2, 1, 1]$, $L_2 = [1, 1, 2]$, $L_3 = [1, 3]$, $L_4 = [3, 1]$ and so on.

Choosing $L_0 = [2, 2, 18, 1, 1]$ with Algorithm 1 gives

$$L_1 = [1, 2, 18, 1, 1]$$

$$L_2 = [2, 18, 1, 1, 2]$$

$$L_3 = [1, 18, 1, 1, 2, 1]$$

$$L_4 = [18, 1, 1, 2, 2]$$

Λ

$$L_{24} = [2, 2, 18, 1, 1].$$

3.3 Terminal Constraints

It is known that stability and feasibility are not ensured by the MPC law without terminal cost or terminal constraints. Usually the problem is augmented with a terminal cost and a terminal constraint set Ψ_f . Typically Ψ_f is a robust control invariant set which guarantees that if

$$\text{Problem } J^\Sigma(x(t), t) = \min_{\hat{u}_{0|t}^\Lambda, \hat{u}_{M-1|t}^s, t_f} \sum_{i=1}^N \{ \mathbb{I}\mathbb{C}(t + i\Delta t) E(x_{i\Delta t|t}, u_{(i-1)\Delta t|t}) \mathbb{I}\mathbb{I}_S + \mathbb{I}\mathbb{I}_{i\Delta t|t} \mathbb{I}\mathbb{I}_u \} \quad (3.2) \text{ is}$$

feasible for a given x_0 , then it is always feasible for $t \geq 0$.

The formal definition for invariant set can be found in [24, 25]. A treatment of sufficient conditions which guarantees persistent feasibility of MPC problems goes beyond the scope of this work and can be found in the surveys [26].

Historical data for $T_{cmp,s}$, $T_{cmp,r}$ and $m_{\&emp}$ are used in order to compute the possible range of \mathcal{Q}_{Load} .

Figure 3.2 Campus Load plots historical daily campus load during September 2008, and it is observed that the load has a period of one day. It is reasonable to model the admissible campus load as a periodic disturbance with periodic envelope constraints (the bounds are represented with thicker lines).

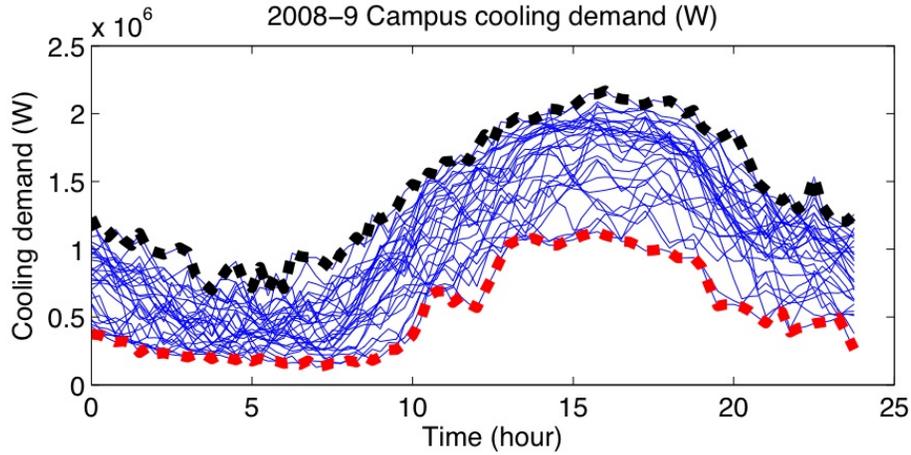


Figure 3.2 Campus Load

Since the disturbance is periodic, the idea proposed by F. Blanchini and W. Ukovich in [27] can be applied to the proposed MPC controller. The invariant sets, if it exists, will be time variant and periodic with the same period as the disturbances. In order to guarantee that the tank has enough cold water to satisfy the demand, the algorithm proposed in [27] is used to calculate the CPI (Controlled Periodic Invariant) set for the system described in Equation 2.20. The system for calculating the CPI set is a simple buffer plant subject to the periodic disturbance shown in Figure 5.6.

The algorithm proposed in [27] was implemented and Figure 3.3 Campus Load Lower bound of the Controlled Periodic Invariant set plots the lower bound $b(t)$ of the computed periodic set $\Psi_f(t)$:

$$\Psi_f(t) = \{z_b(t) : b(t) \leq z_b(t) \leq z_{tank}\} \quad (3.4)$$

If the height of the cooler water in the tank is greater than the lower bounds, $b(t)$, there exists a feedback control law that will satisfy any disturbance belonging to the envelope in Figure 3.3 Campus Load Lower bound of the Controlled Periodic Invariant set without violating the states and inputs constraint.

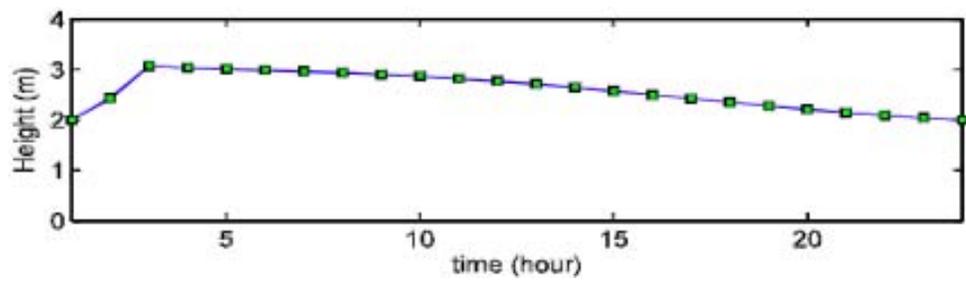


Figure 3.3 Campus Load Lower bound of the Controlled Periodic Invariant set

4 Experiments

The simulation-based estimated performance of predictive algorithms was planned to be demonstrated through an experiment performed in coordination with UC Merced facility staff during the warmer part of the year. During the execution of a three-week experiment in June 2009, referred to as Experiment I in this document, and after data analysis, the team learnt that malfunctions of the plant operation and some of the algorithm constraints affected critically the plant operation and resulted in lower performance than the routine manual operation. The unexpected lower performance and its causes, described in Appendix 2, motivated the team to improve both the model and the optimization algorithm and execute a second experiment.

After Experiment I, the performance-critical aspects were reconsidered and several modifications were made regarding model validation and algorithm improvements. For validating the cooling load model, which is critical for tank capacity estimation and therefore overall performance, another experiment was performed in September 2009 with a slightly higher leaving chiller water temperature. The resulting week-long data set was used to validate the model over a wide range of chilled water supply temperatures and in order to ensure that potential savings with higher temperature set-points are explored by the optimization algorithm. This chilled water supply temperature experiment is described in Appendix 3.

Analysis of Experiment I results also revealed that the length of the tank charging window has a critical effect on plant efficiency, and in consequence the optimization algorithm was improved to include the charging time in the optimization variable set.

Comparison of Experiment I data with regular schedule data also revealed that simultaneous variations in ambient conditions (exogenous inputs) and plant controllable inputs introduced significant uncertainty in the estimate of the performance improvement generated by the predictive algorithm. Instead of using week-by-week comparisons between the experiment weeks, a regression model capturing average performance as a function of ambient and some controllable inputs was developed using historical data collected over several months. Although the regression model was only partially validated, it is expected that comparison against this model will provide a more realistic performance improvement of the predictive algorithm.

The model and control algorithm improvements referred to above were included in the optimization algorithm used in Experiment II in October 2009. Section 4.1 provides a general description and rationale for the experiment. Section 4.2 describes procedural details. Section 4.3 presents the experimental results.

4.1 Experiment Description and Rationale

The chiller plant control was realized through three set-points: chilled water mass flow rate, chilled water leaving water temperature and leaving cooling tower water temperature. To ensure consistency of experimental data, the chilled water plant control was implemented such that the levels of the thermocline in the chilled water tank at the beginning and end of each of the experiment time intervals were approximately equal. The three control set-points generated by the optimization algorithm are the reference values for the local control loops for each respective feedback controller, as illustrated in Figure 4.1.

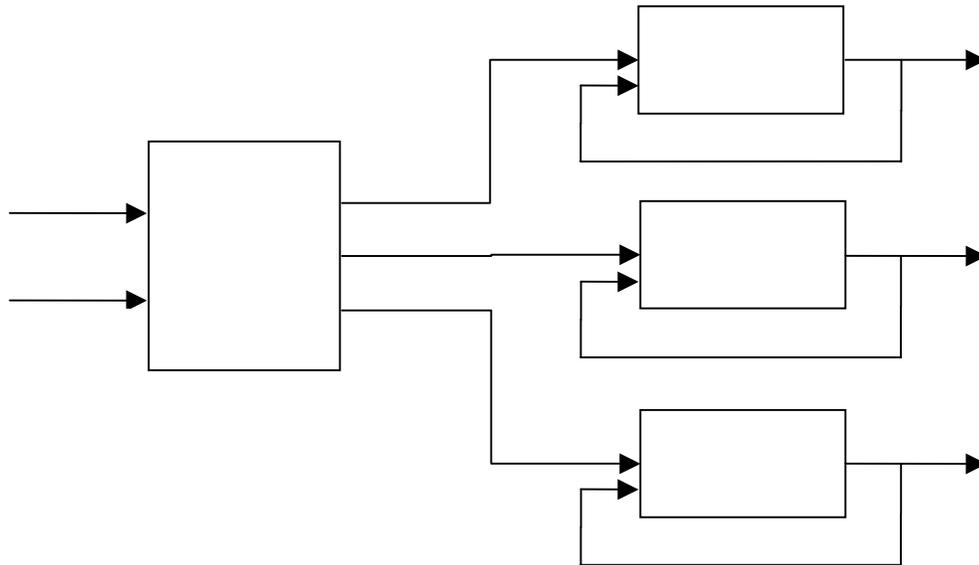


Figure 4.1 Signal flow from MPC algorithm to chiller plant local controllers

The work targeted efficiency improvements in three areas:

- More efficient trade-off between cooling tower and chiller powers, a trade-off that is affected both by the T_{CWS} set-point and the chiller Part-Load Ratio.
- More efficient chiller operation by increasing the chiller PLR through manipulation of the pump flow rate and chilled water leaving temperature set-points, as well as decreased return temperature from the storage tank.
- More efficient storage tank operation by reducing the tank chilled water levels, but sufficient for rejecting predicted thermal loads based on weather forecasts. When^[b10] the chilled water volume stored in the tank exceeds a certain threshold (determined by the thermocline position and height), the temperature of the water above the thermocline decreases and this in turn reduces the chillers Part-Load Ratio, thereby decreasing their efficiencies.

The challenges in the optimization algorithm are a result of the strong coupling between the subsystem efficiencies. As described above, the coupling is a result of the influence of each of the controllable set-points on the performance of the other subsystems.

The overall chillers—thermal storage—campus efficiency for chilled water generation, storage, distribution and consumption depends on both external (exogenous) and controllable inputs and can be written

$$\eta_{Overall} = \eta_{Overall}(T_{CHWS}, T_{CHWS}, \dot{m}_{CHWS}, T_{wb}, T_{CHWR}) \quad (4.1)$$

Controllable inputs
Exogenous inputs

This overall efficiency can be divided into three components as illustrated in Figure 4.2:

- (i) efficiency from generation → storage;
- (ii) storage efficiency, and
- (iii) storage → campus:

$$\eta_{Overall} = \eta_{Plant} \cdot \eta_{Storage} \cdot \eta_{Campus} \quad (4.2)$$

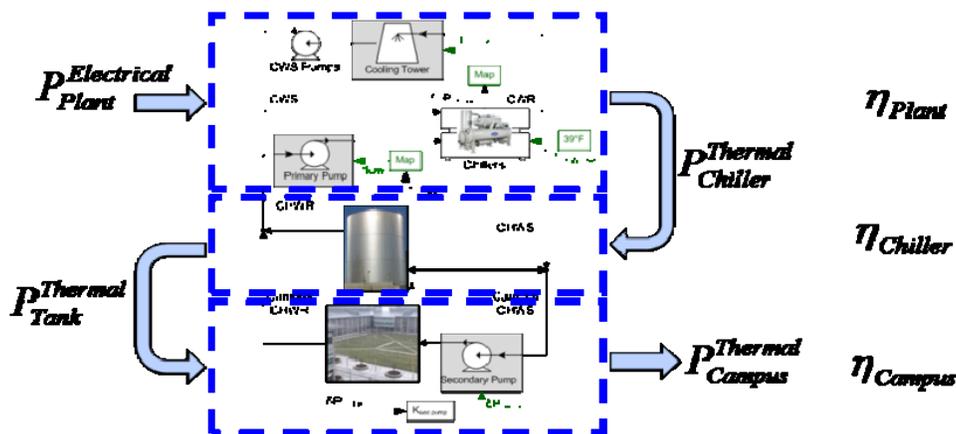


Figure 4.2: Illustration of electrical and thermal power flows for efficiency definitions for chiller, storage tank, and campus

The performance estimation reported here is limited to the first two efficiency components estimation. The third component, η_{Campus} , represents the ratio between the thermal energy delivered to the campus by HVAC system and the thermal energy delivered by thermal storage unit. The only control input, out of the three, that can affect this efficiency is the leaving-chiller water temperature. Although a new value (different from the baseline) of the chiller supply

water temperature set-point may influence the HVAC system efficiency through the campus supply water temperature, this efficiency difference is not accounted for in this report since control of the campus building HVAC system was not within the scope of the work reported here. The second component of the overall efficiency, $\eta_{Storage}$, is computed as the ratio of the thermal energy delivered to the campus, which depends on the mass flow rate to the campus and the campus supply and return temperatures, and the thermal energy delivered by chillers, which depends on the chiller flow rate and the chiller supply and return temperatures. The value of this efficiency was calculated to be close to 1, after the problem with the flow rate measurement had been identified and corrected.

4.2 Experimental Procedures

The second experiment was carried out during the week of October 5-11, 2009. This was not an ideal experimental week, as the cooling load was much lower than during the summer. On the first night (the night of October 5, a holiday), for example, the tank was already nearly fully charged since there was so little load during the preceding day, so the MPC determined that no charging should occur that night (which was the same decision that the standard operator control would have made). Various other problems were encountered during the week ranging from tracking problems to chiller malfunctions. The charging period during the night of October 7 was discarded from consideration in the COP results because of a chiller malfunction, and the charging period during the night of October 11 was also discarded because of another chiller malfunction.

4.2.1 Execution of Experiment II

Experiment II was performed similarly to Experiment I, see description in the Appendix 9.2, and consisted of the sequence of activities of Figure 4.3.

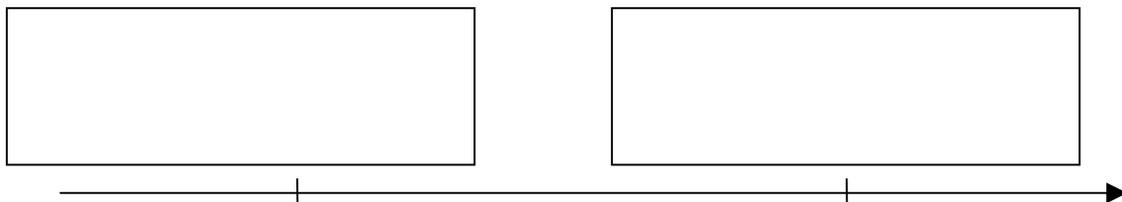


Figure 4.3 Sequence of charging-window activities during MPC experiment II week

The activities mentioned above are detailed below:

1. Data collection: the following data were extracted and used as input for the MPC algorithm

- a. WebCTRL data: stratified storage tank temperatures; chiller return water temperature
- b. NOAA forecast data for the following three days: ambient temperature, humidity.

The data were manually extracted into Excel spreadsheets with specific data structures.

2. Optimal set-points generation: the data collected at the previous step is used as input for the optimization algorithm that generates optimal set-points: $T_{CHWS,ref}, \dot{m}_{CHWS,ref}, T_{CWS,ref}$. These values are optimal set-points computed over a 72h-long predicted horizon.
3. Set-point communication and implementation: the reference set-points computed for the next step were generated as text in the Matlab environment and communicated verbally to the operator who then updated them in ALC interface.

Based on the conclusions of Experiment I described in Appendix 2, two aspects of the experimental procedure and MPC algorithm were improved:

- a) The chiller start time and end time were treated as continuous variables and optimized to achieve the best operation efficiency. (For Experiment I these were set based on operators' experience).
- b) The control sampling time was set to be one day so that the control set-points were kept constant over the charging period of time. (For Experiment I the sampling time was 1 hour.)

The optimal set-points, and chiller start and stop times and were generated once per charging window by running the optimization algorithm at 8:00PM every evening. The choice of the algorithm execution time was made based on the historical time intervals over which the wet bulb temperature has its lowest values (occurring after 8:00PM), and therefore presents the largest potential for high operating efficiency.

4.2.2 Performance Criteria

This section focuses on estimation of the instantaneous (cumulative) plant efficiency, η_{Plant} , defined as the ratio of the thermal power (energy) delivered by chillers to the electrical power (energy); this ratio is known as the Coefficient of Performance and represents the thermal energy generated per kWh of electrical energy consumed:

$$\eta_{Plant}^{Instantaneous} = COP_{Plant}^{Instantaneous} = \frac{p_{Plant}^{Thermal}}{p_{Plant}^{Electrical}} \quad (4.3)$$

$$\eta_{Plant}^{Cumulative} = COP_{Plant}^{Cumulative} = \frac{E_{Plant}^{Thermal}}{E_{Plant}^{Electrical}} \quad (4.4)$$

These performance metrics characterize the MPC policy impact on the chiller plant performance for the following reasons:

- All the optimized set-points affect directly only the central plant's operation efficiency. The three control inputs are coordinated for maximizing the chillers, pumps and cooling towers efficiencies, which can be quantified directly, as illustrated in Figure 2 and above formulae.
- The tertiary HVAC system local control loops are all lumped into a model that predicts the secondary chilled water flow rate, i.e. the flow rate to/from the campus, and the chilled water return temperature.

For calculating the above efficiencies, we make the following assumptions and use the following formulae:

- For $COP_{Plant}^{Cumulative}$ we consider only time intervals for which the initial and final values of the tank capacities are approximately equal. There are two reasons for this constraint. At the beginning and the end of the experiment, the tank level is constrained by the baseline control schedule operation. The second reason is consistency of performance estimation; a difference between initial and final values of the stored water energy can artificially affect the performance estimate, and compensation for this difference would introduce additional uncertainty.

- $$P_{Plant}^{Thermal} = C_P \cdot \dot{m}_{CHW} \cdot (T_{CHWR} - T_{CHWS})$$

- $$P_{Plant}^{Electrical} = P_{CoolingTowers}^{Electrical} + P_{CHWPump}^{Electrical} + P_{PlantChillers}^{Electrical} + P_{ChillerPump}^{Electrical}$$

4.3 Experiment II Data Analysis and MPC Performance Results

Experiment II was executed from October 5th through October 11th, 2009. Compared to Experiment I, executed in June, the lower temperature and cooling load during Experiment II sometimes generated a different tank management strategy. Unfortunately, there were also a few occasions when the chillers and/or cooling towers malfunctioned and the resulted lack of data reduced the statistical significance of the MPC algorithm benefits estimates.

The main control set-points: condenser leaving water temperature, chilled water supply temperature, supply flow rate and chiller start and stop times, are illustrated in Figures 4.4-4.6. With the exceptions of the two nights when the chiller plant malfunctioned, the supply flow rate set-point tracking performance is within the expected accuracy. However, the figures illustrate a larger tracking error for the condenser leaving water temperature set-point, during the nights of October 7th, 8th, and 12th, and for the chilled water supply temperature set-point, during the nights of October 9th – 12th.

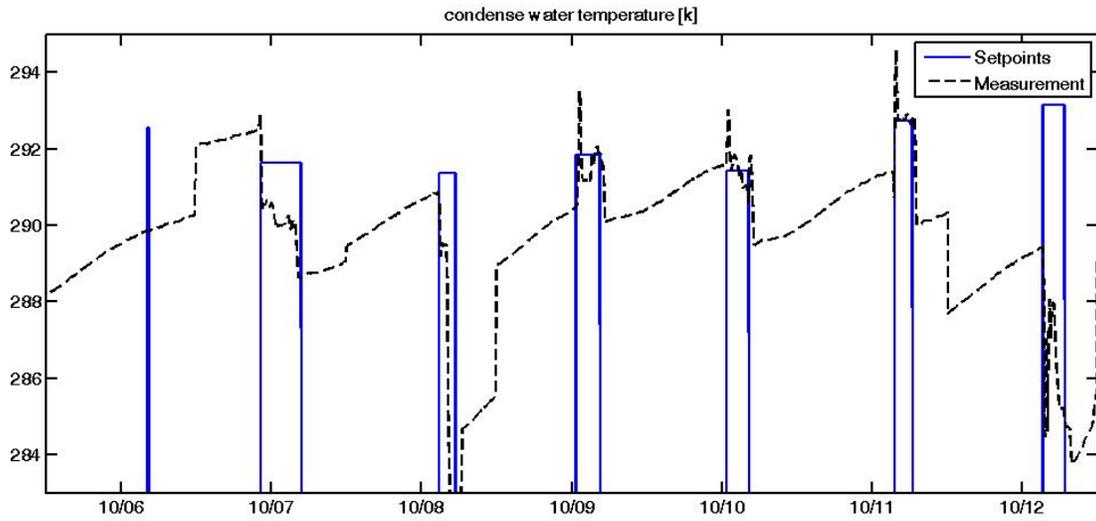


Figure 4.4. Condenser leaving water temperature optimized set-points (solid, blue) and actual values (dashed, black)

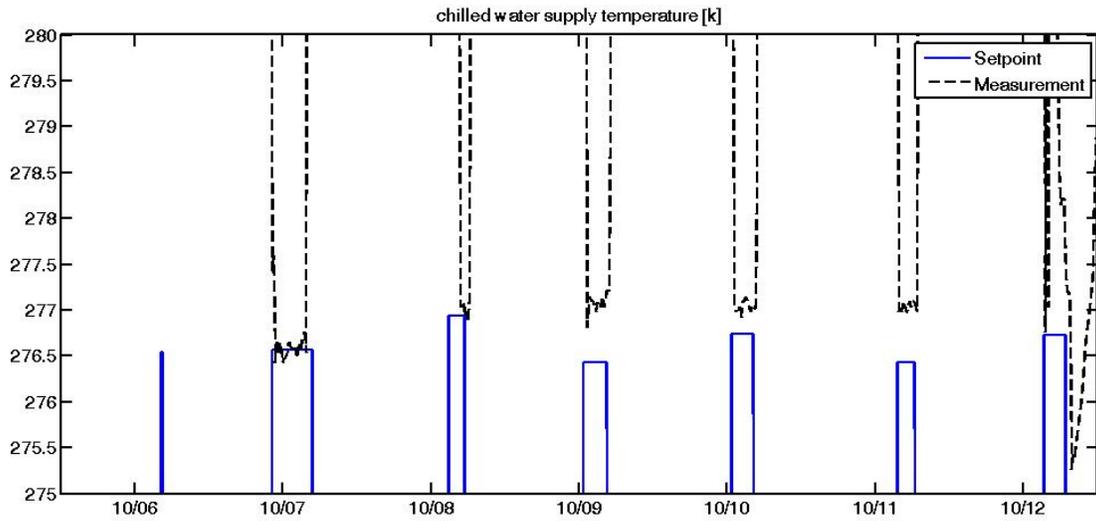


Figure 4.5. Chilled water supply temperature set-point (solid, blue) and actual values (dashed, black)

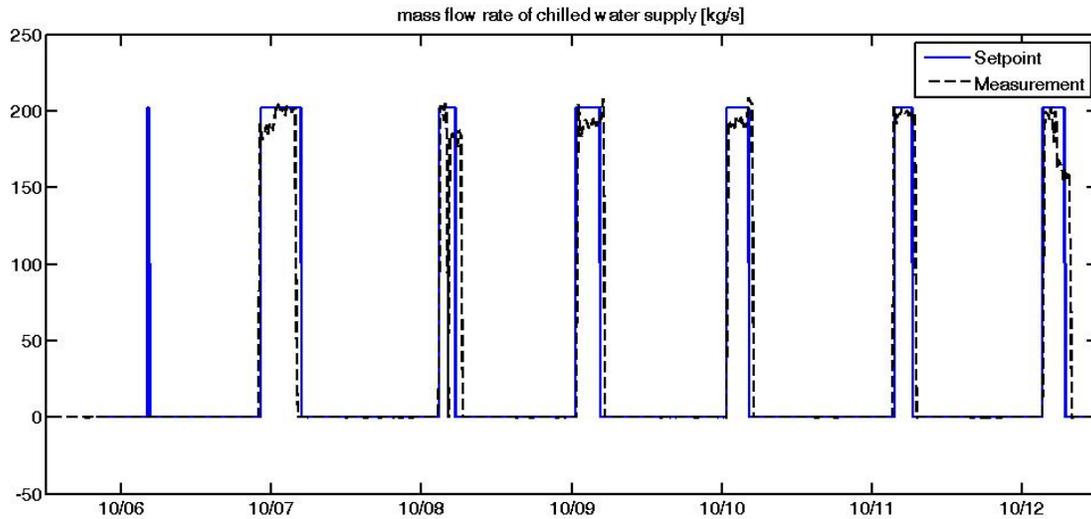


Figure 4.6. Chilled water supply mass flow rate set-point (solid, blue) and actual values (dashed, black)

Figure 4.7 shows the variation in condenser water temperature set-points over different charging periods, as well as the variation in the charging period length specified by the MPC. The details of October 6, 8, 9 and 10, are shown in Figures 4.8, 4.9, 4.10 and 4.11 respectively. The measured average overnight COP for each of the four nights is compared with the two baselines in Figure 4.12.

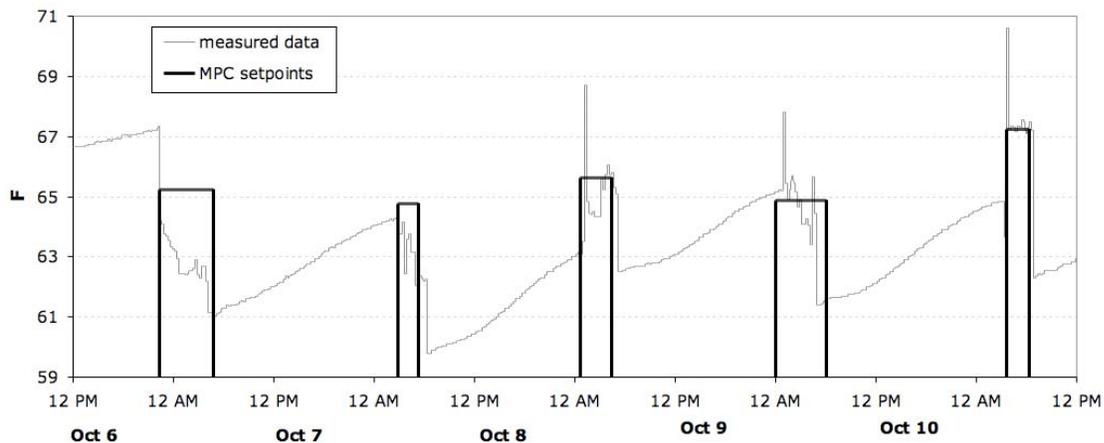


Figure 4.7: Condenser water temperature set-points and measured data, October 6-10 charging periods

The charging period during the night of October 6 was the most successful, in comparison to the two baselines, as shown in

Figure 4.8 below. The charging start time, however, was not much different than the standard start time, and the condenser water temperature was lower than that specified by the MPC and by Baseline #2, which we would expect to be a penalty. The lower than desired flow rate for the first half of the charging period would also be expected to be a COP penalty. However, October 6 was the only day where the chilled water supply temperature tracked our desired set-point of 38°F – this may have increased the COP by increasing the load on the chillers, causing them to operate at a higher COP. It is difficult to draw any conclusions from this single data point, but it is perhaps worth the operators considering decreasing the standard set-point for the chilled water supply temperature.

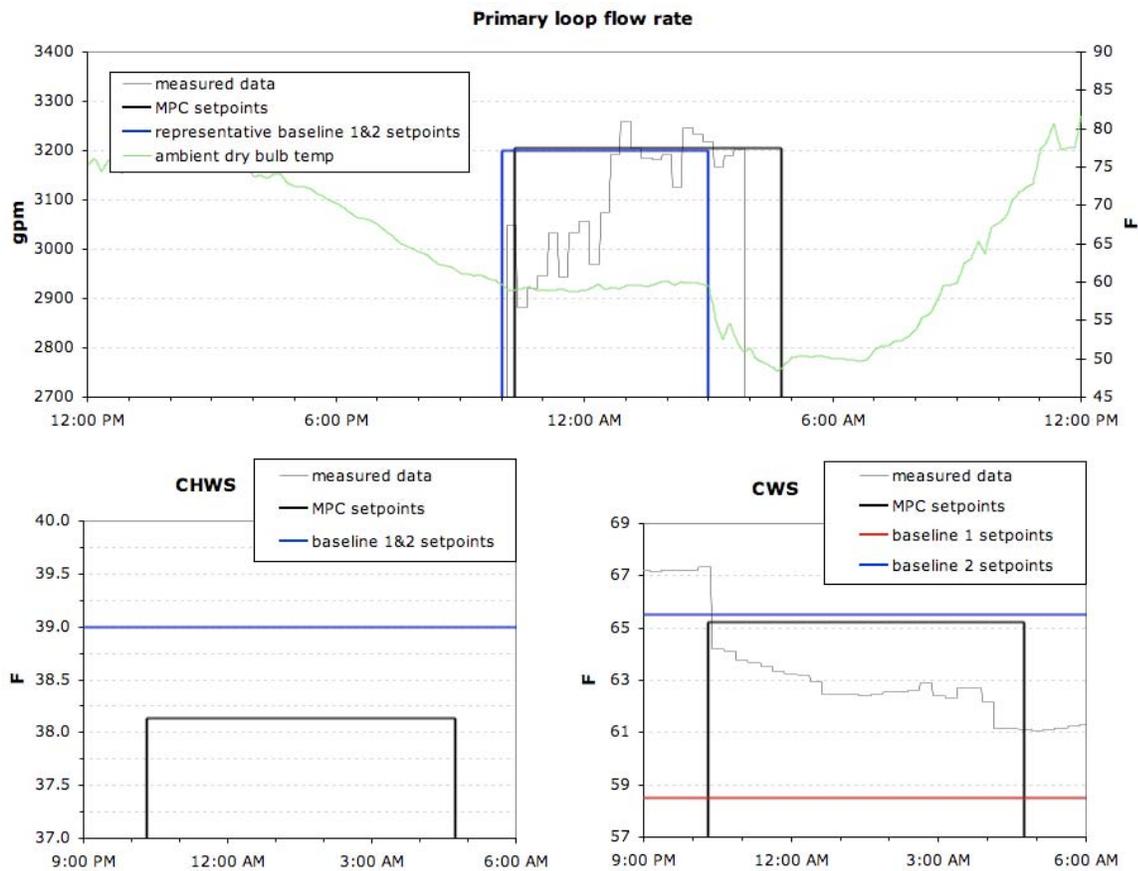


Figure 4.8: Details of MPC operation, charging period on the night of October 6, 2009

During the charging period over the night of October 8, the MPC elected to start charging later in the night than the standard control would have, taking advantage of the lower ambient wet bulb temperature. It is curious, however, that it did not choose to start charging even later in the night, as the ambient temperature continued to decrease. The chilled water supply temperature was tracking the standard set-point rather than the MPC set-point, and the flow rate and condenser water set-point tracking were both rather poor. The measured COP was higher than both the Baseline #1 and Baseline #2 COPs.

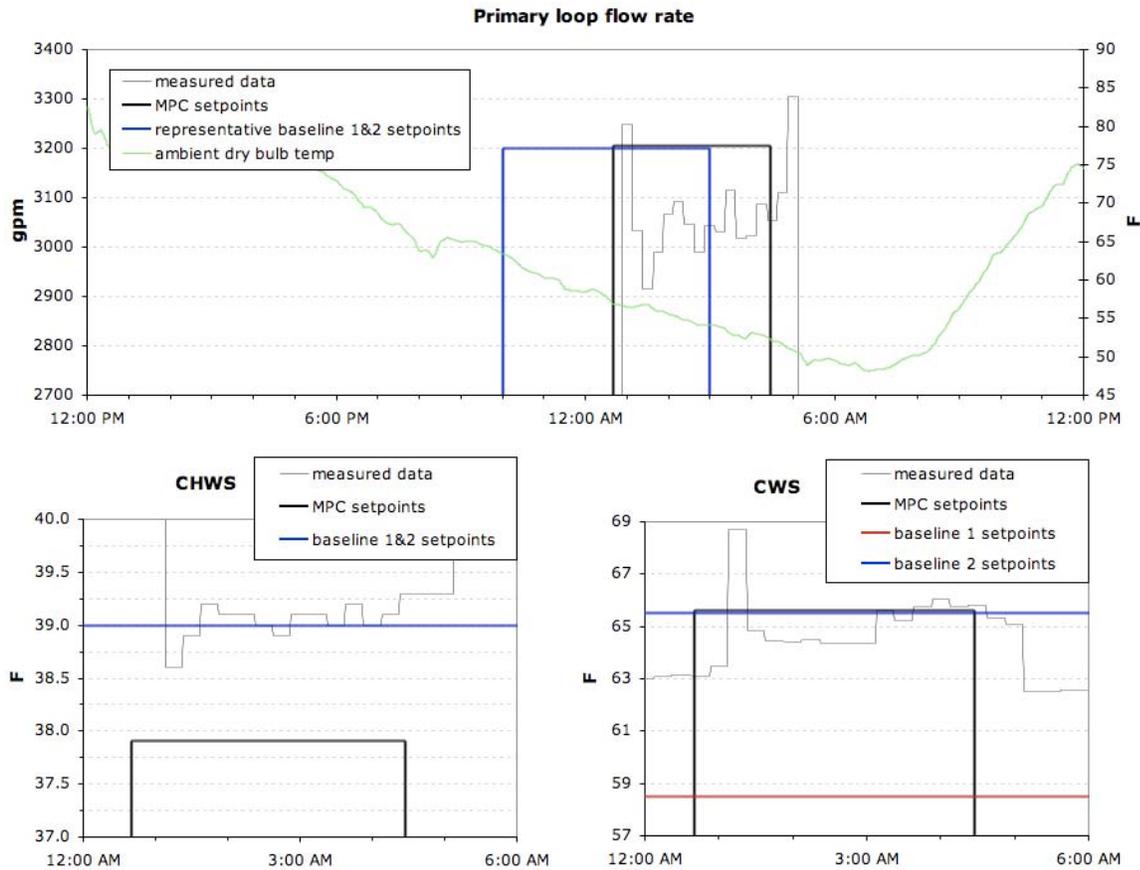


Figure 4.9: Details of MPC operation, charging period on the night of October 8, 2009

The MPC set-points and tracking on the night of October 9 were similar to those on the night of October 8, and the COP results relative to baseline were also similar.

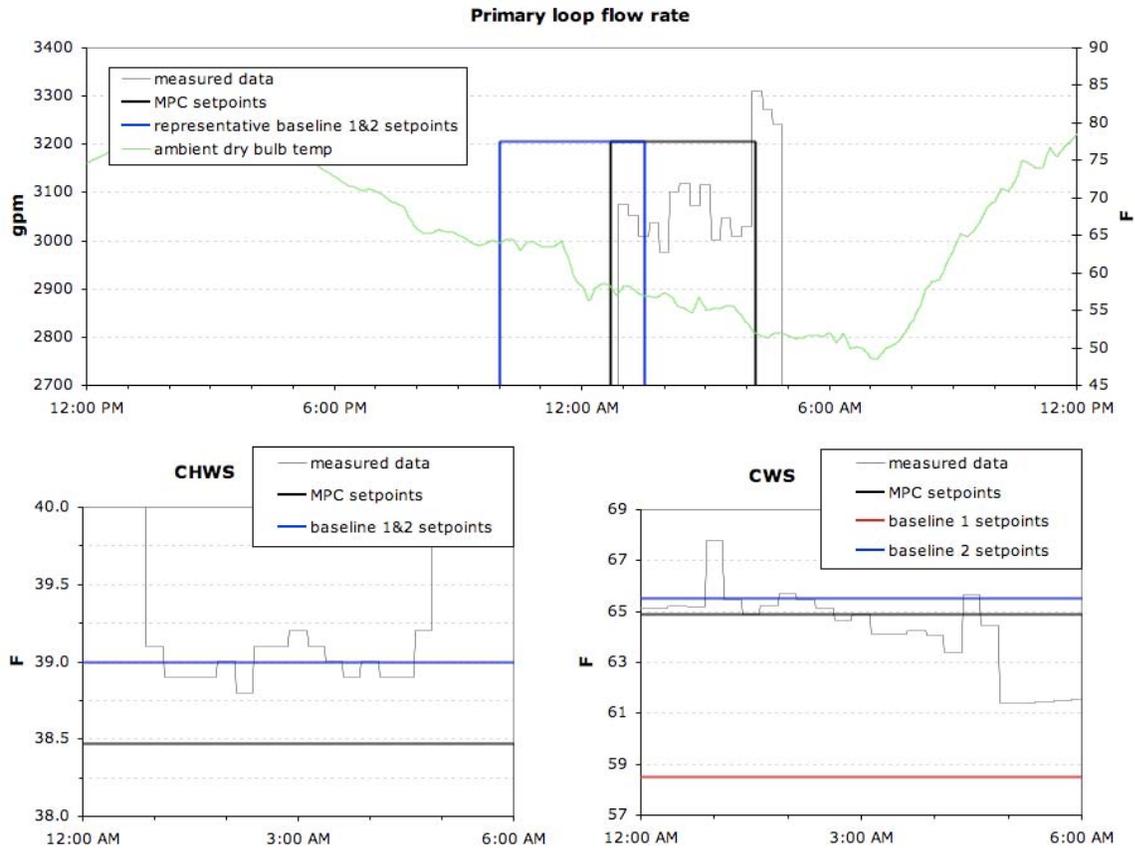


Figure 4.10: Details of MPC operation, charging period on the night of October 9, 2009

The charging period over the night of October 10 seems to have been the most successful of the nights, in terms of the apparent optimality of the charging period (relative to the ambient temperature) and the set-point tracking of the flow rate and condenser water temperature set-point. However, the measured COP for this night was significantly lower than the other MPC nights, and was actually lower than the Baseline #2 COP.

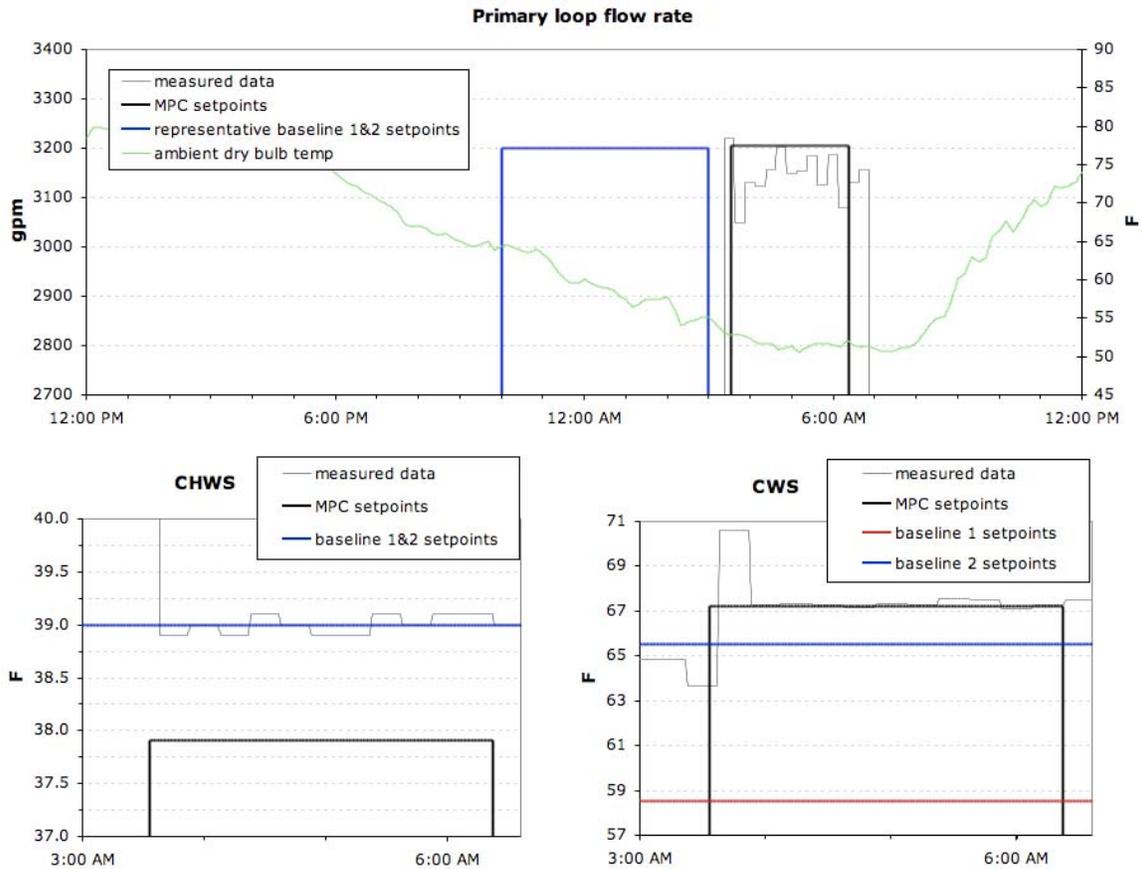


Figure 4.11: Details of MPC operation, charging period on the night of October 10, 2009

Figure 4.12 compares the measured COP during the experiment with the two baselines. The average measured COP was $3.2 \pm 2.2\%$ higher than Baseline #2 (where 2.2% is one standard deviation in the accuracy of the baseline regression model), and $4.6 \pm 2.4\%$ higher than Baseline #1.

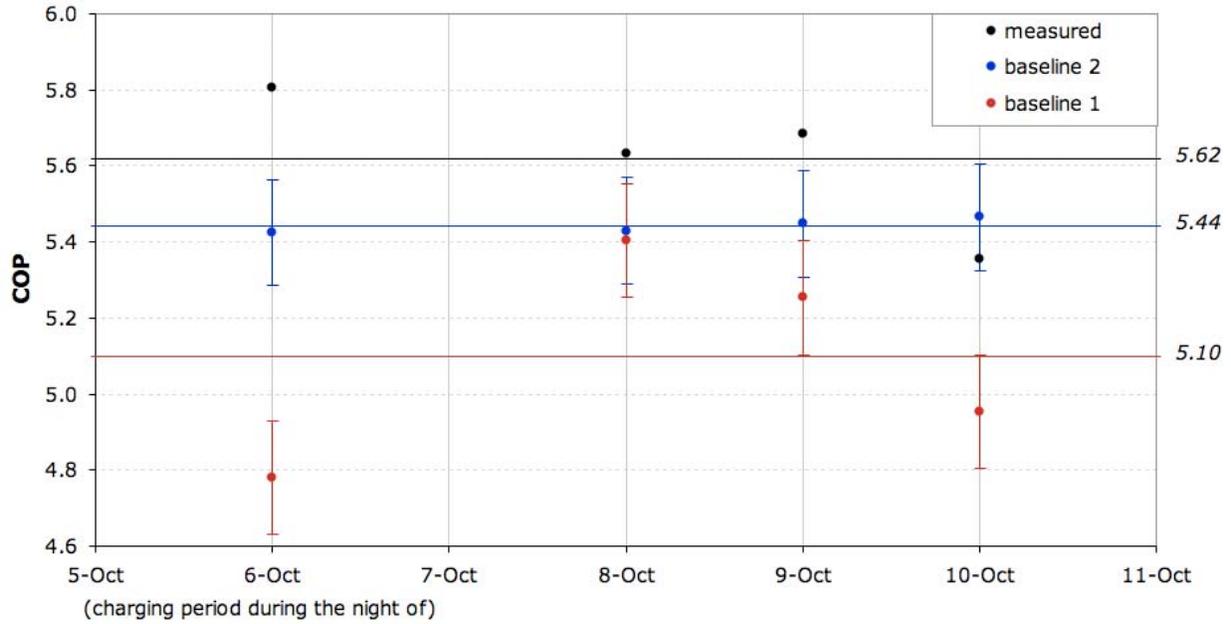


Figure 4.12: Average charging period COP, measured vs. baselines

5 Discussion

The process of data collection and model development, necessary for any MPC project, resulted in the team uncovering various problems with the chilled water system. This is discussed in Section 5.1: Chilled Water Plant Operational Problems Identified. Although it is difficult to quantify the energy savings resulting from these diagnostics, they were likely on the same order as the energy savings from the MPC itself. And although the types of problems uncovered and the amount of energy savings may differ significantly from project to project, some benefits of diagnostics are expected with the use of MPC for any chilled water plant.

Although saving energy at the UC Merced campus is a worthwhile goal in itself, this research project was intended as a demonstration of what is possible for any such system. This is discussed in Section 5.2: Applicability of Results.

The final subsection, Section 5.3: Challenges for Commercial Feasibility, notes implementation challenges that were encountered during the project along with a general discussion of challenges such as model calibration and controller robustness that will need to be addressed in order for MPC to see significant market penetration.

5.1 Chilled Water Plant Operational Problems Identified

5.1.1 Primary Loop Flow-Rate Sensor

In the process of identifying the parameters of the chiller models, the chilled water flow rate sensor was determined to be unreliable. The chiller cooling load, the rate at which heat is removed from the primary chilled water loop by the chiller, is calculated as $\dot{Q}_{load} = C_p \dot{m}_{chw} (T_{chwr} - T_{chws})$. Using an energy balance on the chiller condenser, the quantity of heat rejected to the condenser loop of the chiller, \dot{Q}_{cond} , should satisfy $\dot{Q}_{cond} = \dot{Q}_{load} + P_{ch}$. Given \dot{Q}_{cond} and P_{ch} , the condenser energy balance yields an estimate of the chiller cooling load $\hat{\dot{Q}}_{load} = \dot{Q}_{cond} - P$. As shown in Figure 5.2, the cooling load measured for Chiller #3 calculated via the flow-rate measurement agrees poorly with the cooling load estimated using the condenser energy balance. In contrast, the cooling load measured for Chiller #2 matches the cooling load estimated using the condenser energy balance, see Figure 5.3. The flow-rate sensor is placed just downstream of a point where the flow from Chiller #3 emerges from a right angle bend whereas the flow from Chiller #2 is straight for a sufficient number of pipe diameters to provide a reasonable measurement point. Thus, it was determined that the flow-rate sensor would be significantly more accurate when Chiller #2 was the second Chiller, as opposed to Chiller #3. Consequently, the experiments were only executed with Chiller #2 as the second chiller.

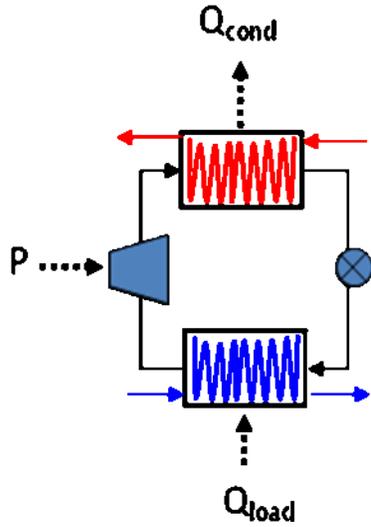


Figure 5.1: Power and heat flow diagram for chiller

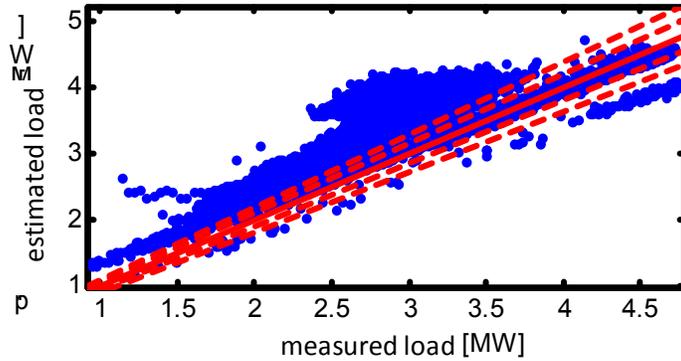


Figure 5.2: Chiller #3 comparison of estimated cooling load versus measured cooling load

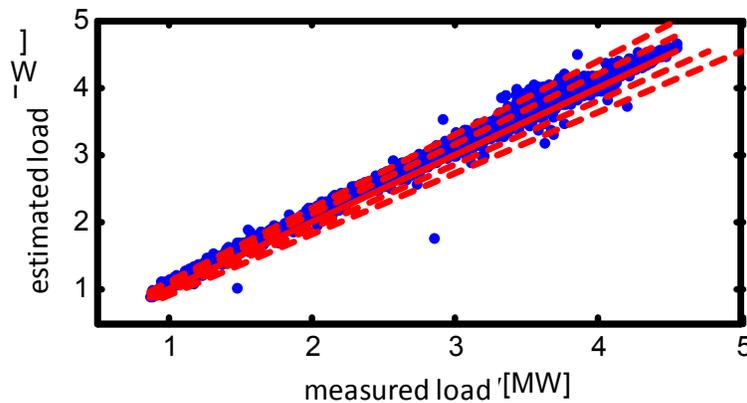


Figure 5.3: Chiller #2 comparison of estimated cooling load versus measured cooling load

5.1.2 Secondary and Tertiary Loop Flow-Rate Sensors

The secondary loop flow-rate sensor located at chilled water plant was determined to be completely unreliable due to its placement relative to pipe fittings; therefore, the sum of the flow-rate sensors located in each tertiary (building) loop was used to calculate the secondary loop flow-rate. However, it was found the tertiary loop flow-rate measurements were also suspect immediately prior to the first experiment in June 2009.

The first check was based on an energy balance. That is, the chilled water return temperature coming from the campus was compared with the mass weighted chilled water return temperatures of the tertiary loops. The chilled water return temperature coming from the campus was significantly lower than the chilled water return temperature predicted by the mass-weighted value, thereby suggesting

that one or more flow-rate meters were reading incorrectly.

Figure 5.4 compares the weighted chilled water return temperatures of the tertiary loops with the centrally measured campus return temperature, alongside a temperature measurement from near the top of the tank, which corroborates the central return temperature measurement.

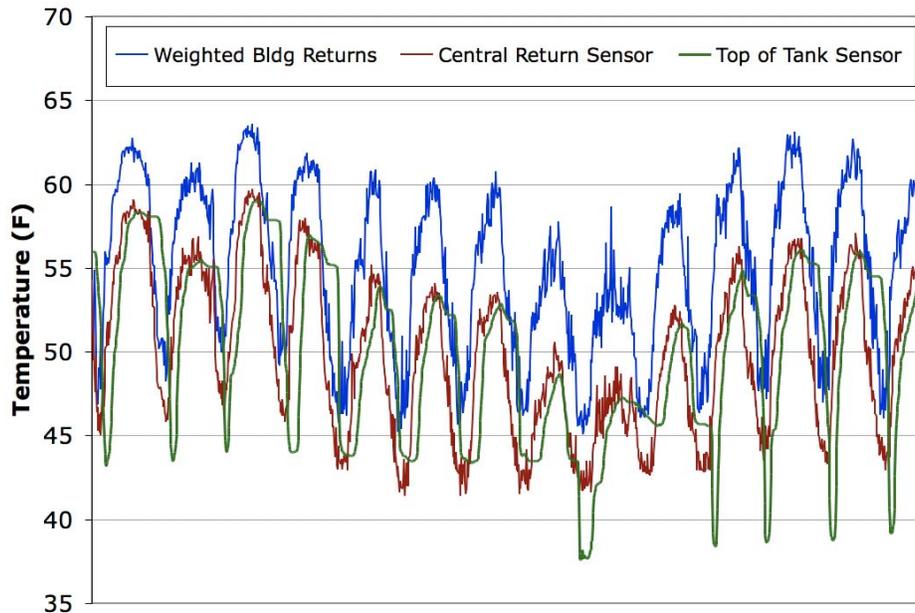


Figure 5.4: Comparison of return temperatures measured centrally, in the tank, and the return temperatures from the buildings weighted by their measured flow rates

Figure 5.5 compares the daily load on the chillers to the daily campus load as calculated at the buildings level and the daily campus load as calculated centrally using the measured return temperature and the sum of the measured flow rates. This analysis strengthened the suggestion that some of the chilled water flow was not being accounted for, that the missing flow rate was significant, and that the campus return temperature was being underestimated as a consequence. This problem was not only making the modeling task difficult, since it was difficult to know which data to trust in calibrating models, but also likely decreasing the efficiency of the chilled water system itself.

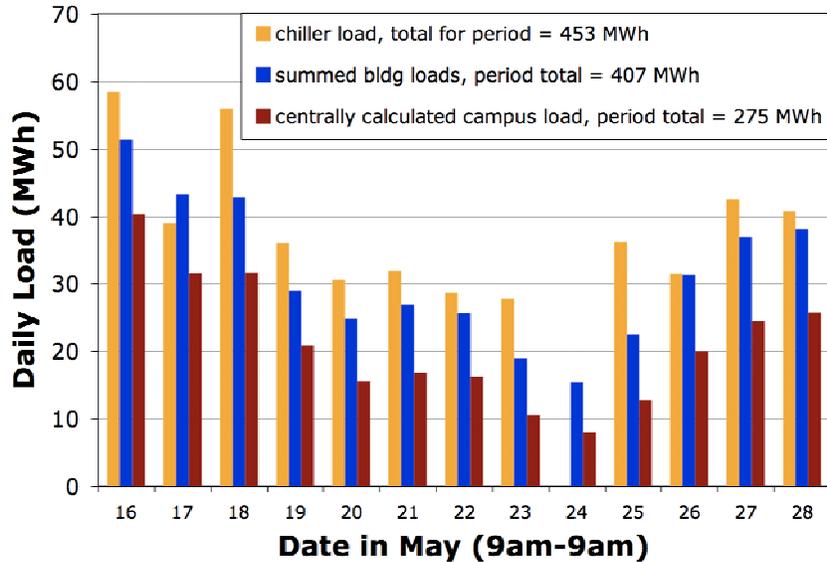


Figure 5.5: Comparison of the daily chiller loads, the daily loads calculated at each building and daily campus load calculated centrally

The second check was based on a mass balance. The position of the thermocline in the tank was used as a means of estimating the flow-rate. As shown in Figure 5.6, the flow-rate sensor consistently under-predicted the flow rate of chilled water from the thermal energy storage tank. As a result, the mass balance confirms that the tertiary loop flow-rate measurements underestimated the total secondary flow-rate.

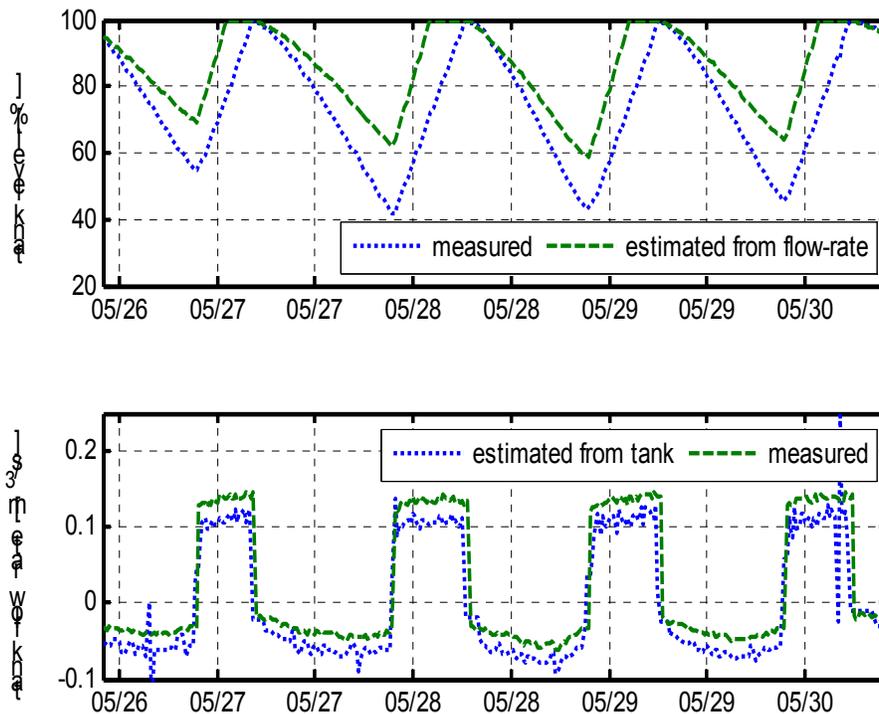


Figure 5.6: Comparison of measured and estimated tank levels

The problem was remedied during June 2009 when the UC Merced HVAC technicians recalibrated several suspect tertiary flow meters. The fix increased the campus return temperature, thus increasing the central system’s COP, and also significantly decreased the pump power consumption at one of the buildings, where the pump was being controlled with a faulty sensor and was thus running at full capacity most of the time.

5.1.3 Chiller Surging Prediction

An energy balance [PH11] was used to provide the corrected flow rate used for the chiller calibration. The corrected flow-rate measurements and calibrated chiller maps were used to analyze chiller surging phenomena reported by UC-Merced. Figure 5.7 shows that when Chiller #2 was unable to maintain its 38°F [PH12] chilled water supply set-point, the adjusted part load ratio (PLR) predicted by the model was greater than 1. That is, the model is able to predict when the duty demanded of the chiller exceeds its capacity.

The UC Merced central plant manager noted that chiller surging had been observed to be a problem during particularly hot summer conditions. In fact, hot ambient conditions do not completely characterize the occurrence of chiller surging. Chiller surging depends on the chiller loading, which depends on the cooling demand, evaporator temperature and condenser temperature, rather than only one of the boundary conditions. Currently, the leaving chiller temperature set-point is kept constant and the pump flow-rate is regulated to maintain an approximately constant cooling load as the chilled water return temperature varies. It is hypothesized that chiller surging occurs whenever the demand on the chiller is greater than its capacity at the current condenser and evaporator temperatures:

$$\dot{m}_{chw} C_p (T_{chwr} - T_{chws,ref}) > Q_{cap}(T_{cws}, T_{chws}). \quad (5.1)$$

This quantity may also be captured in terms of the adjusted part load

$$PLR_{adj} = \dot{m}_{chw} C_p (T_{chwr} - T_{chws,ref}) / Q_{cap} < 1. \quad (5.2)$$

Figure 5.7: Chiller #2 surging captured by $PLR_{adj} > 1$, start day is Aug. 20, 2008; (blue chw = chwr.ch2; green chw = chws.ch2; blue cw = cwr.ch2; green cw = cws.ch2)

Figure 5.7 shows more than 60 days of the summer of 2008., The surging condition resulted in Chiller #2 being unable to maintain its 38°F set-point. The times when Chiller #2 was unable to maintain its set-point coincides with the times when $PLR_{adj} > 1$. Figure 5.7 strongly suggests that $PLR_{adj} > 1$ can be used to predict chiller surging conditions.

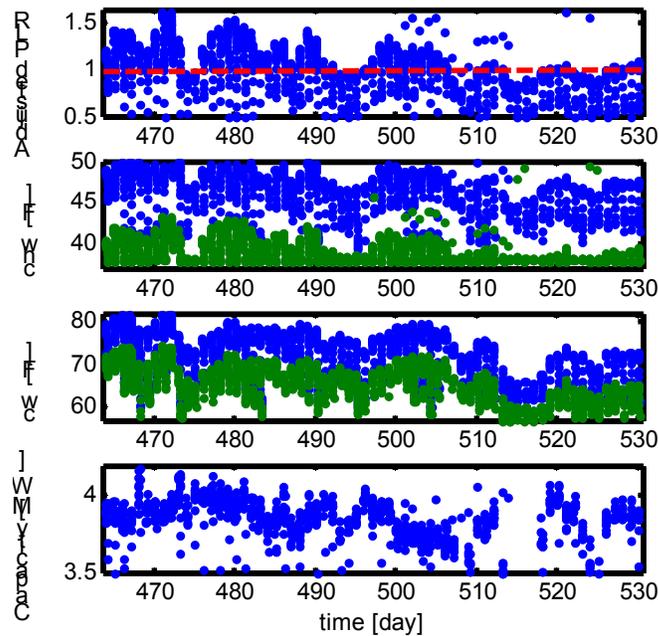


Figure 5.7: Chiller #2 surging captured by $PLRadj > 1$, start day is Aug. 20, 2008; (blue chw = chw.r.ch2; green chw = chws.ch2; blue cw = cwr.ch2; green cw = cws.ch2)

5.1.4 Ambient Humidity Sensor

It was observed the ambient humidity sensor located at the University of California Merced indicated very small variations in relative humidity. The weather records for the Merced Municipal Airport weather station were downloaded for comparison. Although the dry bulb temperature records for Merced Municipal Airport closely agree with the UC Merced records, the Merced Municipal Airport records indicated significant variation in the relative humidity, see Figure 5.8. As shown in Figure 5.9, the cooling tower model calibrated with the suspect relative humidity sensor was of poor quality. Consequently, the airport measurements were used in place of the campus measurements for calibrating the cooling tower model, and the issue was brought to the attention of the UC Merced facility staff. As shown in Figure 2.23, the quality of the model fit was significantly improved by using the airport relative humidity sensor.

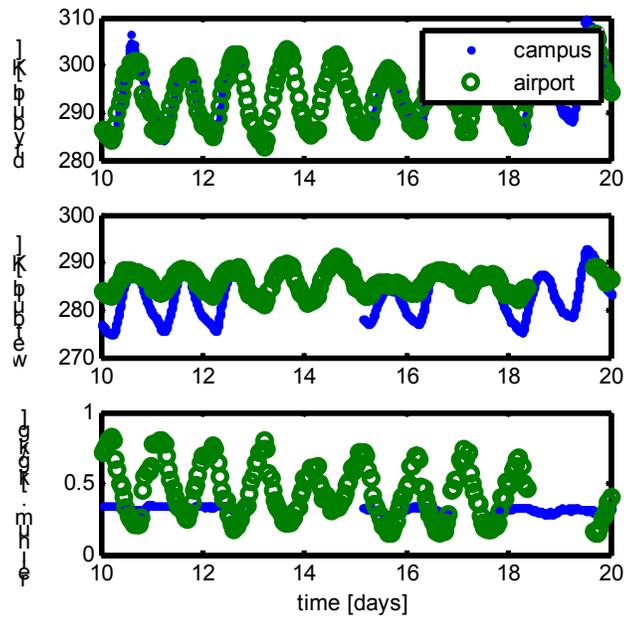


Figure 5.8: Comparison of campus weather station to airport weather station measurements

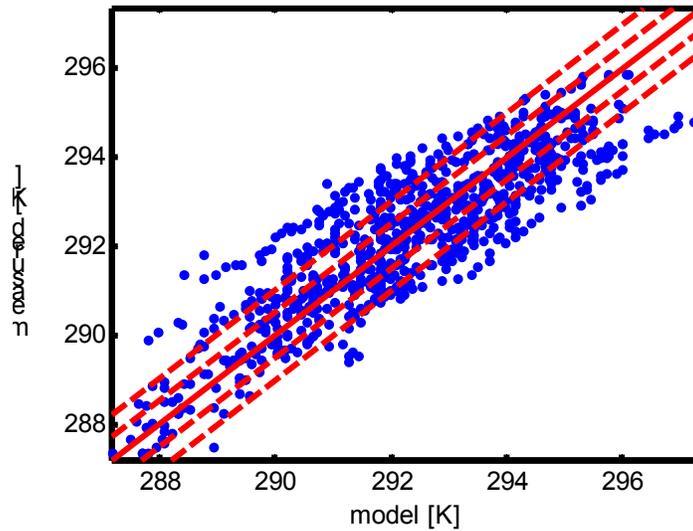


Figure 5.9: Comparison of predicted and modeled condenser water supply temperature using model calibrated from UC Merced relative humidity sensor.

5.1.5 Degraded Cooling Tower Performance

The measured performance of the cooling towers was compared to data from the manufacturer, Baltimore Aircoil Company, who provided the performance curves shown in

Figure 5.10. Data from the cooling towers at UC Merced are plotted against the manufacturer's curves in Figure 5.11. Based on the data, the leaving cooling tower temperature is

approximately 2.5 K (4.5°F) higher than predicted by the manufacturer’s curves. Besides the actual hardware not meeting specifications, this discrepancy could be caused by a number of installation-related problems. For example, the cooling towers are located on top of the chilled water plant and architectural screening is used to hide the equipment. The screening and other factors may lead to the recirculation of the high humidity air exiting the cooling towers back into the cooling towers. The reason for the substantially lower cooling tower performance is under investigation in a separate project, and the impact on overall chilled water plant power consumption will be quantified.

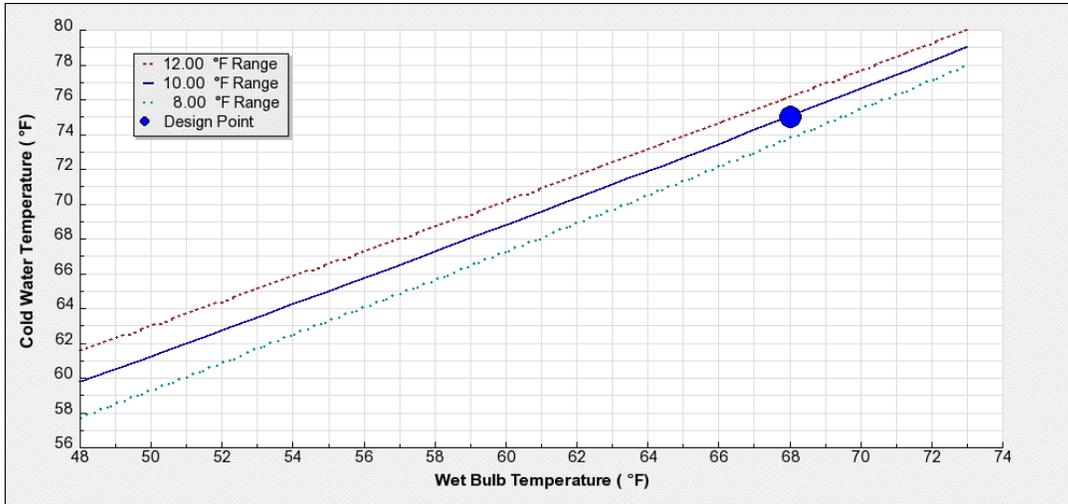


Figure 5.10: Cooling tower performance curves provided by Baltimore Aircoil Company

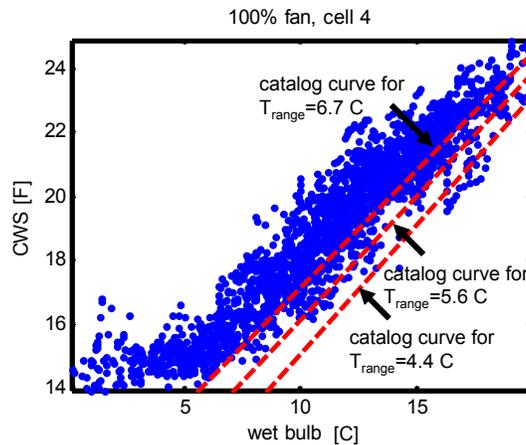


Figure 5.11: Comparison of measured and predicted leaving temperature

5.1.6 Chilled Water Storage Tank Operation

In addition to the commissioning benefits discussed above, the MPC development and testing process was also able to identify simple ways to improve the heuristic control policy currently used by the operators. It was found that operating the chillers near full load was a key factor in maximizing system efficiency. In order to maximize the chiller load, one must maximize the temperature difference across chillers and the chilled water flow rate through the chillers. However, the admissible chilled water supply temperature and flow rates are bounded, so the chiller loading is constrained and often determined by the temperature of the primary loop chilled water return, which is primarily a function of the temperature near the top of the TES tank. This temperature falls substantially as the thermocline approaches the top of the tank, reducing the available chiller loading. Thus, overcharging the storage tank can be detrimental to the chilled water plant efficiency and should be avoided. It is also expected that additional savings are available by concentrating low loads on a single chiller rather than spreading it across multiple chillers. As noted earlier, the plant operators learned from the MPC set-points in this first experiment that the system COP can be increased by increasing the standard condenser water set-point range. Further research is required to determine how much energy could be saved through the addition of simple rules like these to the existing policy, based on the lessons learned in the MPC study.

5.2 Applicability of Results

The UC Merced campus is currently only a fraction of its future planned size and the storage tank is sized for a substantially larger campus cooling load than currently occurs. In particular, the capacity of the tank is always sufficient to meet the cooling load for the next day and there is no need to operate the chillers during the day, which simplifies the control problem. In order to assess the performance of the MPC approach implemented in the project under more typical conditions, a follow-on simulation study will be performed to investigate the expected performance as the load from the campus increases.

5.3 Challenges for Commercial Feasibility

Many challenges remain for commercial feasibility of model predictive control of chilled water plants with thermal energy storage. Further development of the MPC algorithms is only a small aspect to the challenges requiring a comprehensive platform that also includes modeling, calibration, controller interface, and fault detection and diagnostics.

For similarly configured chilled water plants it is expected only minor modifications would be required to employ the same modeling approach detailed above. However, the calibration process was found to be quite labor intensive. First, it was not possible to rely on design specifications and manufacturer specifications, because the performance of many chilled water plant components failed to meet design specifications, e.g. the primary pump flow-rate and the

cooling towers. Second, care is needed to ensure that historical data used to calibrate models cover the operating envelope of the component in question. Third, the performance of components was subject to change after maintenance and failures. For example, the primary pump balancing valve adjustment affected the system curve and it was discovered prior to the first MPC experiment one of the cooling tower cells was inoperable.

Additional measures can be taken to make the model predictive control algorithm more robust to such model discrepancies. Such measures take the form of worst case analysis to ensure the chilled water plant requirements are met for the entire set of model uncertainties of concern. Nevertheless, the overall chilled water plant efficiency will undoubtedly suffer if the model is not immediately adjusted for such discrepancies. Consequently, a model-based fault detection and diagnostics is one potential way of detecting and reacting to discrepancies in the chilled water plant performance. Many of the faults discovered in the process of executing the project were 'soft' faults that are not detected by the alarm capabilities of current control systems and hence did not halt the operation of the chilled water plant. However, these faults significantly degraded the COP of the chilled water plant and compromised any subsequent attempts to optimize its efficiency. Furthermore, current building control systems tend to inundate the operators with alarms, many of which are false alarms. As result, most faults go unheeded by the operators given that the operators are without a reliable fault detection mechanism and a way of prioritizing the significances of faults

6 Conclusions and Recommendations

A Model Predictive Control (MPC) scheme was successfully developed and implemented for the UC Merced chilled water system. The energy savings relative to the previous manual operation were $4.6 \pm 2.4\%$. The operators modified their manual operation procedures in the light of the results of the first MPC test and the MPC savings relative to this new baseline were $3.1 \pm 2.2\%$.

Two experimental tests of the MPC were carried out. The first test revealed flaws in the initial implementation, which were subsequently fixed. However, as a result of the first experiment, the system operators learned that a higher condenser water temperature set-point would improve performance and changed their operating procedures accordingly, complicating the establishment of a baseline. After the first experiment, a primary pump was also fixed, which increased the system COP slightly. A faulty flow sensor in one of the buildings was also found as a result of our MPC development and implementation (the 'bypass' problem), and was fixed just after the first experiment, also causing an increase in COP because of the higher return temperatures. A second MPC experiment (October 2009) was then carried out with the corrected MPC controller.

Two different regression-based baselines were used for comparison of the primary system COP during the second MPC experiment against the COP under normal operation. "Baseline 1" is based on the system performance before the first experiment, and thus estimates what the COP would have been if there had been no intervention at all by the MPC research team. "Baseline 2" is based on the system performance between the first experiment and the second experiment, and thus represents good system operation with near-optimal conventional control with static set-points. The measured COP during the MPC experiment was higher than Baseline 1 by $4.6\% \pm 2.7\%$, and higher than Baseline 2 by $3.1\% \pm 2.2\%$.

In the course of the MPC project, various problems with the chilled water system operation were discovered, largely as a part of modeling process. The flow-rate measurements were the most susceptible to significant measurement errors. The first inaccurate flow meter was located on the primary chilled water loop and the second inaccurate flow meter was found on one of the tertiary (building) chilled water loops. The latter also revealed a fault in one of the building air handling units. It was determined the cooling towers failed to meet manufacturer specifications. Prior to the experiments, it was found that one of the interconnects for a cooling tower cell fan had failed, and the fan had been inoperable for weeks if not months. Many of the problems are typified as a soft fault. That is, the problems did not cause the system to fail, and there were no overt signs in the status of what is a complex system.

6.1 Future Work

The UC Merced campus is currently only a fraction of its future planned size and the TES tank is sized for a substantially larger campus cooling load. Consequently, an upcoming extension of the research is to further characterize via simulations the advantages of MPC scheme to chilled water systems similar to UC Merced with different storage capacity, cooling loads and climate.

The above-described MPC implementation is not robust in the sense that neither input uncertainties nor model error are taken into account. This could be addressed by designing a robust MPC scheme [19, 20, 27] (Borrelli 2003; Witsenhausen 1968; Bemporad, Borrelli & Morari 2003). A typical robust MPC scheme involves solving a min-max problem to optimize robust performance (the minimum over the control input of the maximum over the disturbance) while enforcing input and state constraints for all possible admissible bounded disturbances. Also, the MPC implementation does not consider system faults. Robustness to faults can be indirectly obtained by detuning the controller (reducing the weights) at the price of a reduced performance or can be systematically and easily taken into accounts by switching to a different model $g(x,u,t,d)$ and constraints X and U when a fault is identified. Given the uncertainties and system faults apparent in this case study, and their prevalence in building systems in general, future research on robust MPC for buildings applications is recommended.

In order for the approach to be commercially viable, the time and expertise intensity required for the development and testing of the MPC implementation must be greatly reduced. Ultimately, a streamlined tool chain for the development and implementation of model-based building control and commissioning is needed to address part of these challenges. As discussed, MPC development and fault diagnostics can be two mutually reinforcing activities. In order to facilitate such activities, the development and dissemination of controls-oriented models for building system components, such as those developed during this project, is a first step. The complexity of building systems calls for methods that enable detailed component models to be aggregated into lower-order models for coordination at higher levels. In addition, the configuration of advanced control and monitoring techniques should ultimately be transparent to the installers and operators in order to facilitate communication and understanding and to save energy.

A first step towards a commercially viable implementation is described in Appendix 4. The approach will be tested in the summer of 2010, taking advantage of the repairs and enhancements to the instrumentation of the chilled water plant and distribution system implemented by the UC Merced Facilities staff in the first half of 2010. As discussed above, problems with chilled water flow rate measurements, in particular, limited both the performance and the assessment of the MPC algorithm in 2009.

7 Acknowledgements

The authors wish to thank the following for their invaluable assistance: John Elliott and Scott Walling – University of California Merced, Scott Bortof and Satish Narayanan – United Technologies Research Center and Michael Wetter – Lawrence Berkeley National Laboratory.

This work was supported by the Assistant Secretary for Energy Efficiency and Renewable Energy, Office of Building Technology, State and Community Programs of the U.S. Department of Energy under Contract No. DE-AC02-05CH11231 and by the California Energy Commission PIER Buildings program through the California Institute for Energy and the Environment.

8 References

- [1] Morari, M. and J. Lee. 1999. "Model predictive control: past, present and future," *Computers and Chemical Engineering* 23 667–682.
- [2] Mayne, D., J. Rawlings, C. Rao, and P. Scokaert. 2000. "Constrained model predictive control: stability and optimality," *Automatica* 36 (6) 789–814.
- [3] Borrelli, F., A. Bemporad, M. Fodor, and D. Hrovat. 2006. "An MPC/hybrid system approach to traction control." *IEEE Transactions of Control Systems and Technology* 14 (3)
- [4] TIAX, LLC. Energy impact of commercial building controls and performance diagnostics: Market characterization, energy impact of building faults and energy savings potential. Report No. D0180, Prepared for the DOE Building Technologies Program. 2005.
- [5] CBECS 2003. Commercial Buildings Energy Consumption Survey commercial energy uses and costs. Energy Information Administration
- [6] B. Erpelding, "Real Efficiencies of Central Plants," *HPAC Engineering*, May 2007
- [7] Henze, G.P., M. Krarti, and M.J. Brandemuehl. "Guidelines for improved performance of ice storage systems." *Energy and Buildings*, 35(2):111 -- 127, 2003.
- [8] Brambley, M.R., P. Haves, S.C. McDonald, P. Torcellini, D. Hansen, D.R. Holmberg, K.W. Roth, *Advanced Sensors and Controls for Building Applications: Market Assessment and Potential R&D Pathways*. Pacific Northwest National Laboratory, NL-15149, 2005.
- [9] Wetter, W. and Polak, E. "Building design optimization using a convergent pattern search algorithm with adaptive precision simulations." *Energy and Buildings*, 37(6):603– 612, 2005.
- [10] Wetter, M. and Wright, J.A. "A comparison of deterministic and probabilistic optimization algorithms for nonsmooth simulation-based optimization." *Building and Environment*, 39(8):989–999, 2004.
- [11] Grottschel, M., Krumke, S. O. and Rambau, J. Eds., "Online Optimization of Large Scale Systems," Springer, 2001.
- [12] Coffey, B., Haghghat F, Morofsky E, and Kutrowski E. 2010. "A Software Framework for Model Predictive Control with GenOpt." *Energy and Buildings*, in press
- [13] Wetter, M.. 2009. "Modelica-based Modeling and Simulation to Support Research and Development in Building Energy and Control Systems." *Journal of Building Performance Simulation*, Vol. 2, No. 2, pp. 143-161.
- [14] Wetter, M.. 2009. "Modelica-based Modeling and Simulation to Support Research and Development in Building Energy and Control Systems." In *Proc. of the 11th IBPSA Conference*, pp. 652--659. Glasgow, Scotland, July.
- [15] <http://www.dymola.com>

- [16] <http://www.mathworks.com>
- [17] Energy Plus Engineering Reference Manual, April 2008.
- [18] Hydeman, M., Webb, N., Sreedharan, P., and Blanc, S., "Development and Testing of a Reformulated Regression-Based Electric Chiller Model," ASHRAE Transactions, v. 108, no. 2, 2002.
- [19] Holmes, M.J., The simulation of heating and cooling coils for performance analysis", Proceedings of System Simulation in Buildings '82, Liège, Belgium, December, 1982.
- [20] F. Borrelli, 2003. "Constrained Optimal Control of Linear and Hybrid Systems". Lecture Notes in Control and Information Sciences, Springer-Verlag. Vol 290.
- [21] H.S. Witsenhausen. 1968. "A min-max control problem for sampled linear systems". IEEE Trans. Automatic Control, 13(1):5–21.
- [22] S. Qin and T. Badgwell. 1997. "An overview of industrial model predictive control Technology," Control Engineering Practice,
- [23] R. Cagienard, P. Grieder, E. Kerrigan, and M. Morari. 2007. "Move blocking strategies in receding horizon control," Journal of Process Control, vol. 17, no. 6, pp. 563 – 570.
- [24] F. Blanchini. 1999. "Set invariance in control — a survey," Automatica, vol. 35, no. 11, pp. 1747–1768, Nov.
- [25] A. Bemporad and M. Morari, 1999, "Robust model predictive control: A survey," in Robustness in Identification and Control. Springer-Verlag, no. 245, pp. 207–226.
- [26] D. Mayne, J. Rawlings, C. Rao, and P. Scokaert, 2000. "Constrained model predictive control: Stability and optimality," Automatic, vol. 36, no. 6, pp. 789–814, Jun.
- [27] F. Blanchini and W. Ukovich, 1993. "Linear programming approach to the control of discrete-time periodic systems with uncertain inputs," J. Optim. Theory Appl., vol. 78, no. 3, pp. 523–539.
- [28] Bemporad, A., F. Borrelli, and M. Morari. 2003. "Min-max Control of Constrained Uncertain Discrete-Time Linear Systems," IEEE Transaction on Automatic Control, Vol. 48, No. 9, September, pp. 1600- 1606.

9 Appendices

9.1 Appendix 1: Identification of a Baseline Performance Model

The objective is to quantify the improvement in the Coefficient of Performance (COP) for the chilled water system obtained with Model Predictive Controller (MPC) as compared to the standard way of operating the system. This is not a straightforward task. One possible approach is to use the model of the chilled water system that was developed for the project driven by the heuristic control logics used by the operators prior to the project. However, we can reduce the error if we use the following approach, which is to identify a baseline COP prediction model by fitting a fourth order polynomial to measured data using linear regression

$$COP = Poly(\bar{m}_{CHWS}, \bar{T}_{CHWR,cmp}, \bar{T}_{wb}) = \sum_{i=0}^4 \sum_{j=0}^{4-i} \sum_{k=0}^{4-i-j} \alpha_{ijk} \bar{m}_{CHWS}^i \bar{T}_{CHWR,cmp}^j \bar{T}_{wb}^k \quad (9.1)$$

where, \bar{m}_{CHWS} is the average mass flow rate produced by the chillers during the charging time (kg/s), $\bar{T}_{CHWR,cmp}$ is the mass-weighted average of the campus chilled water return temperature for one day preceding start of charging (F), and \bar{T}_{wb} is the average overnight wet bulb temperature during charging (F). There are two sets of data to be fitted. Note that these models can only be used for validation purposes.

- Baseline1 (B1) is from Dec. 2nd, 2008 to Jun. 1st, 2009. The central plant is operated manually according the policy generated by plant managers.
- Baseline2 (B2) is from Jun. 7th, 2009 to Oct. 6th, 2009. This is the period of time following the first MPC experiment. The central plant operators took advantage of Experiment I by reducing the set-points for the condenser water supply temperature.

We denote the regression model fitted by data from B1 and B2 as model COP1 and model COP2, respectively. The coefficients of COP1 and COP2 models are listed in Table 9.3.

Table 9.3. Coefficient of Performance regression model coefficients

	α_{000}	α_{100}	α_{010}	α_{001}	α_{200}	α_{110}	α_{101}	α_{020}
B1	5.9666	1.7624	6.3583	-8.3794	0.2222	-0.6516	-0.8134	1.0281
B2	1.1227	13.5040	15.0881	14.6085	0.3785	-4.2457	-0.4187	3.4655
	α_{300}	α_{210}	α_{201}	α_{120}	α_{111}	α_{102}	α_{003}	α_{021}
B1	-0.0018	0.0022	0.0059	0.0027	0.0062	-0.0069	-0.0126	-0.0115
B2	-0.0016	0.0089	-0.0024	0.0717	-0.0181	0.0215	-0.0502	-0.1988
	α_{400}	α_{310}	α_{301}	α_{220}	α_{211}	α_{202}	α_{130}	α_{121}

B1	4.83e-6	-1.06e-5	-1.81e-5	1.99e-5	1.97e-5	2.27e-5	-3.14e-5	-9.30e-5
B2	4.53e-6	-2.69e-5	-1.11e-5	-2.84e-5	1.37e-4	1.40e-5	-5.74e-5	3.10e-4
	α_{040}	α_{031}	α_{022}	α_{013}	α_{004}	α_{011}	α_{002}	α_{012}
B1	5.56e-5	1.29e-4	7.22e-5	4.66e-5	1.86e-5	0.4905	1.3608	-0.0091
B2	-2.80e-6	0.0026	-0.0020	0.0012	-1.15e-4	5.2524	-1.4676	0.0925
	α_{112}	α_{103}	α_{003}					
B1	-3.03e-5	3.83e-6	-0.0070					
B2	-5.23e-4	-2.21e-6	-0.0321					

Figure 9.12 presents the results of the linear regression. The dots are the measured COP and the crosses are COP predicted by the regression model.

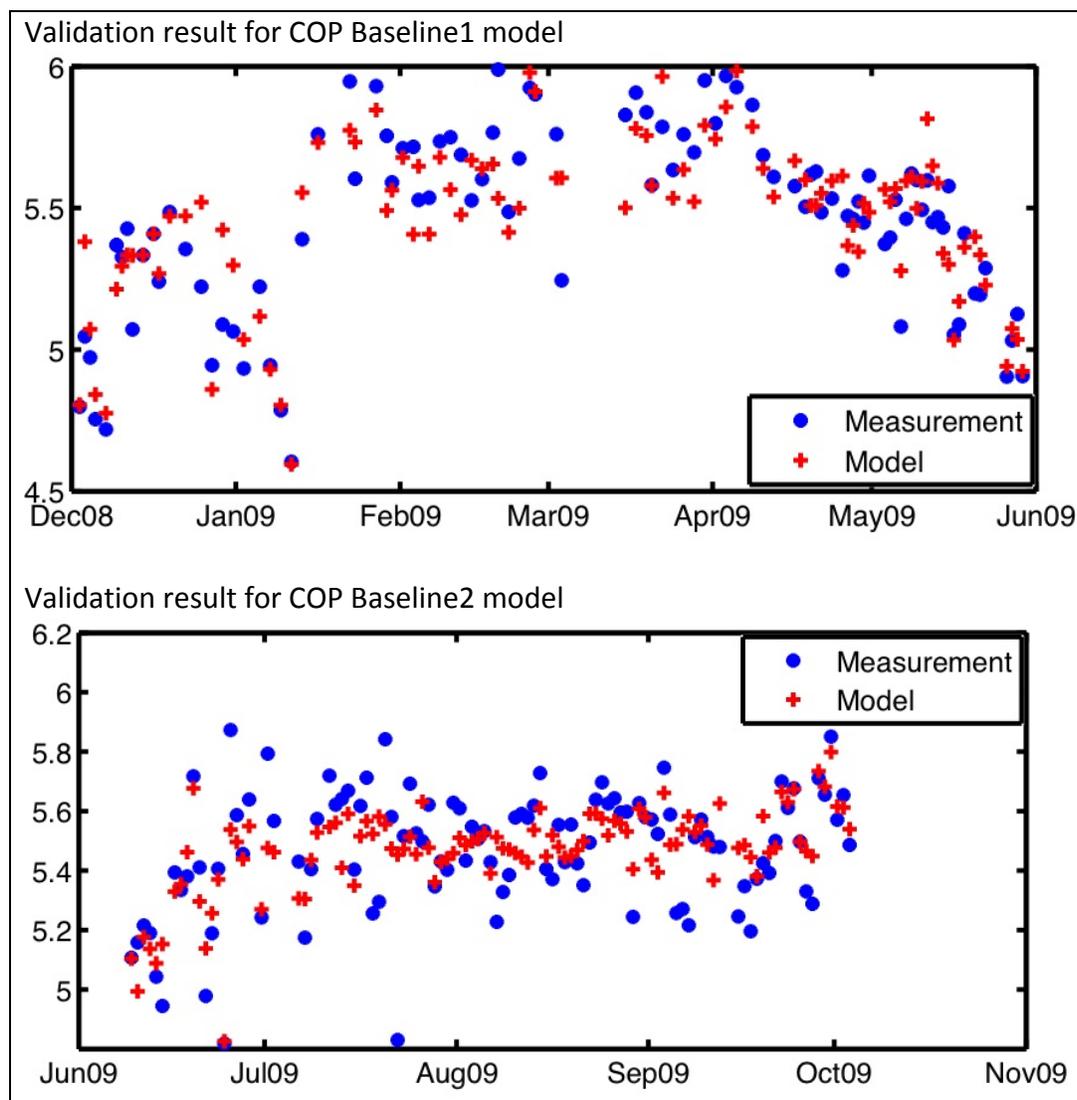


Figure 9.12. COP linear regression validation results

The absolute errors of the COP1 and COP2 regression models are plotted in Figure 9.13, and a statistical characterization of the model errors is presented in Table 9.4, which shows that the prediction COP regression models COP1 and COP2 have r.m.s. errors of 0.15 and 0.14, respectively.

Table 9.4. Error statistics for linear regression models

Error Statistics	Max	Min	Standard Deviation	Mean [$\times 10^{-6}$]
<i>B1</i>	<i>0.45</i>	<i>-0.36</i>	<i>0.15</i>	<i>3.86</i>
<i>B2</i>	<i>0.33</i>	<i>-0.62</i>	<i>0.14</i>	<i>-1.61</i>

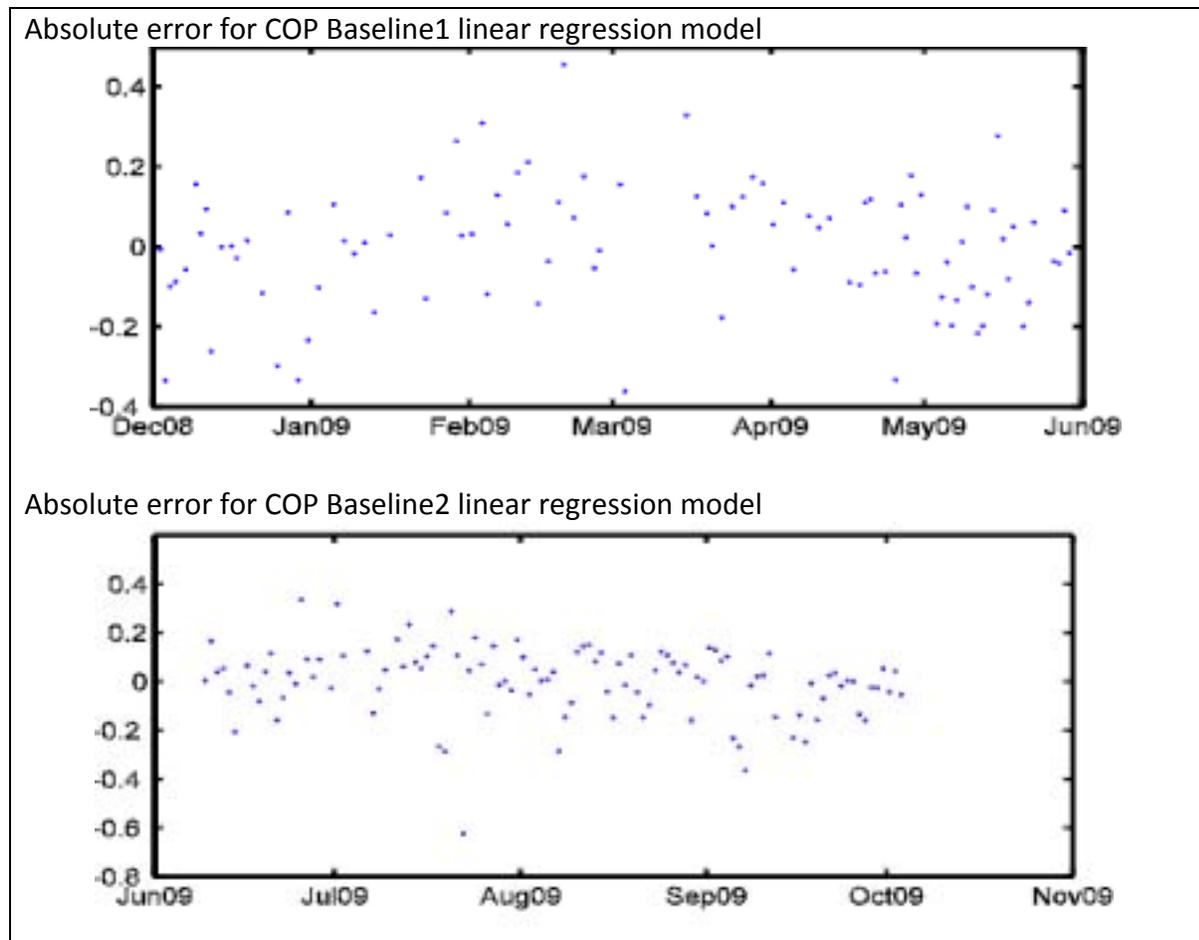


Figure 9.13. Absolute error of COP linear regression^[PH13]

9.2 Appendix 2: MPC Experiment I

For validating the simulation-based estimated energy savings, Experiment I was conducted in coordination with UC Merced facility staff over a three-week time interval from Sunday, May 24th, 9:00PM, through Sunday, June 14th, 9:00PM. For ensuring consistent ambient conditions and occupancy patterns, and thus generating consistent experimental results and meaningful analysis, the plant was operated with baseline set-points—schedules that operator use on a daily basis—during the first and third experiment weeks. The Model-Based Predictive Control (MPC) algorithm generated set-points that were implemented during the second week. Although efforts were made to ensure the desired consistency, changes in ambient conditions, practical aspects of the plant operation, and occasional lack of coordination between the MPC and the plant technical teams have resulted in temporarily operational deviations from the planned three-week experiment schedule.

9.2.1 Execution of Experiment I

9.2.1.1 MPC-Based Set-points Generation and Implementation Procedure

During the MPC policy experimentation week (second experiment week), the plant set-points were generated on line, with an one hour sampling interval, during the storage tank charging window using the optimization algorithm with updated actual measurements from The building control system and weather forecast from NOAA server. Once computed, the set-points were communicated to the operator who updated them accordingly using the ALC interface. Figure 9.29 illustrates the sequence of activities during the tank-charging windows of the MPC week.

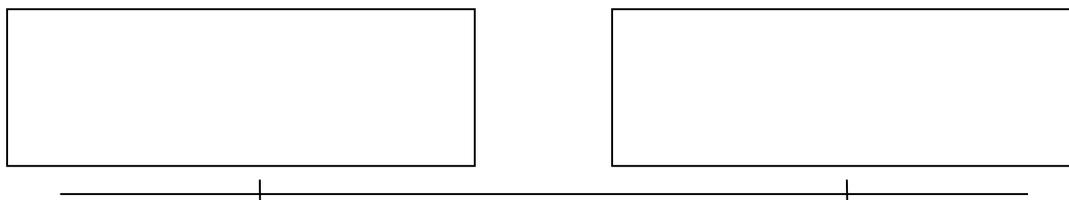


Figure 9.29. Sequence of charging-window activities during MPC experiment week

The activities mentioned above are detailed below:

1. Data collection: the following data is extracted and used as input for the MPC algorithm
 - a. WebCTRL data: stratified storage tank temperatures; chiller return water temperature
 - b. NOAA forecast data for the following three days: ambient temperature, humidity.

The data was manually extracted into Excel spreadsheets with specific data

- structure. These files were updated accordingly into the optimization algorithm folder.
2. Optimal set-points generation: the data collected at the previous step is used as input for the optimization algorithm that generates optimal set-points: $T_{CHWS,ref}$, $\dot{m}_{CHWS,ref}$, $T_{CWS,ref}$. These values are optimal set-points computed over a 24h-long predicted horizon.
 3. Set-point communication and implementation: the reference set-points computed for the next step were generated as text in the Matlab environment and communicated verbally to the operator who then updated them in ALC interface.

9.2.1.2 Baseline Week 1

The first (out of three) week of the experiment was used as a baseline for estimating energy saving benefits. In addition to ensuring ambient conditions consistency with the MPC week, the data and information collected on-site during this week were used to:

- Integrate the overall on-line algorithm consisting of all components and subsystem models, weather forecast data extraction function, cooling load prediction
- Test off-line the newly integrated algorithm and estimate its performance in a simulation environment.
- Perform additional model re-calibration and validation
- Investigate optimization algorithm performance sensitivity with respect to algorithm parameters

Prior to Week 1, earlier versions of the optimization algorithm were used to provide initial estimates of performance improvements; these versions used off-line load estimation. During Week 1 all the algorithms for data extraction, filtering, load estimation, and optimization were integrated for automatic generation of set-points based on input data files. The integration consisted in implementing the data transfer between various subroutines generated with various tools: Excel, Dymola, Matlab, and Simulink, and ensuring the consistency of input/output variables, type and format, and of assumptions used by all subroutines.

The re-calibration activities were focused on updating the cooling load models based on the out-of-season occupancy pattern and on what initially appeared to be, and later was confirmed by the energy manager to be, either a short circuit in the campus chiller water loop or a malfunctioning flow meter. It was later confirmed, in the third experiment week, that there was a short circuit in the water loop caused by a failure in a building recently added to campus chilled water loop (Common Terrace building). For the MPC week, the coil and campus load models were based on a bypass model whose accuracy is illustrated in earlier sections of this report.

The baseline week 1 measurement data were also used for offline study of the optimization algorithm performance and for fine tuning of several parameters: sampling rate, number of optimization steps, charging window length. The performance estimation for the baseline

schedule during this week was affected by the deviation from normal chiller plant operation due to two events:

- Chiller #1 surged unexpectedly and had to be restarted manually. During the 0.5 hours long re-start process there was no chiller water flow and the efficiency decreased. The chiller's surge cause was not identified.
- Cooling tower cell did not operate for a few nights and had to be repaired.
- Pumps operated at flow rates below typical values and this resulted in smaller efficiency and an unrealistic performance reference value. The pump capacity was restored during the MPC policy week (Week 2).

The performance estimation for baseline Week 1 did not compensate for these two anomalies.

9.2.1.3 Model Predictive Policy (Experiment Week 2)

The prediction-based algorithm was implemented during the 2nd week following the steps mentioned briefly in Section 9.2.1.1. These activities are detailed in the following script used by the team for coordination during this week:

1. 1PM : Daily policy preview for upcoming night
 - a. Extract data from ALC for initialization (tank state and cooling load model)
 - b. Run MPC algorithm to generate preview of policy for upcoming night
 - i. Policy generates information and plots for following step
 - c. MPC Team: review and iterate on policy
 - i. Analyze control sequence and compare to previous policies.
 - ii. Constraint violation: chiller, cooling tower, tank, campus pump flow-rate
 - iii. Break down of energy usage
 - iv. Review reduced-order model tank states
 - d. Archive data and policy for post-analysis
2. 3PM : Team briefing with plant operations
 - a. Previous night's performance
 - b. Policy preview
 - c. Gather recommendations and concerns from plant operators
 - d. Implement critical mitigation actions
3. 5PM : Send out daily update
4. 8PM – 6AM : Hourly Implementation Procedure
 - a. hh:00 Collect data from ALC to initialize MPC
 - b. hh:10 Generate policy
 - c. hh:30 Check policy via predefined metrics. Constraints are satisfied. Is there a large gap when compared to policy preview? If policy fails check, implement backup policy.
 - d. hh:45 Archive data from current step
 - e. hh:60 Operator implements new set-points

5. 9am : Daily post-analysis
 - a. Extract data from previous day from ALC
 - b. Collect MPC input generated over night
 - c. Simulate model response with data from ALC and implemented MPC
 - i. Compare to preview and actual policy
 - ii. Identify constraint violation: chiller, cooling tower, tank, campus pump flow-rate
 - iii. Compare actual and simulated chilled water plant performance via break down of energy usage
 - iv. Compare actual and simulated tank states
 - d. Archive data and policy for post-analysis
6. Mid-week review: course correction
 - a. Review model predictions versus actual performance (from ALC)
 - b. Verify admissible process variables: tank level, temperature, power break down, chiller operation, cooling load met (critical AHU supply air temperatures, tertiary pump speeds, secondary pump speeds, fan power)

Two main operational aspects, with an unknown influence on final performance, deviated from the initially planned MPC experiment: testing of the boiler performance and a gradual change in condenser set-point constraint. The boiler test temporarily increased the cooling load for a maximum of one hour. It is expected that this load was relatively small compared to the campus daily load, and its effect on overall performance is within normal operational variability. This latter effect is addressed in the “Analysis of Experiment” section. The overall performance was affected significantly by the incremental change in the leaving condenser water set-point. From the beginning through the middle of the testing week, the upper bound constraint on the condenser leaving water set-point was gradually increased to the optimized value originally recommended by the MPC policy. This was a conservative approach for avoiding chiller surging under new operational constraints. However, the influence of the new set-point on the overall performance was difficult to estimate (reference the “Analysis of Experiment” section) due to oscillatory response of the local condenser control loop. During the third week this behavior was corrected in the Energy Management Control System (EMCS).

9.2.1.4 Baseline Week 3

To ensure consistency in the ambient conditions, the week following implementation of the MPC policy was considered part of Experiment I. However, the condenser water set-point was not reset to its regular operation (prior to MPC policy week) and this improved the reference performance value used as baseline. In addition, the re-tuning of the condenser PID gains improved the tracking of the set-point, which is generated by the optimization algorithm, illustrating to a larger extent than in Week 2 the influence of an increase in the leaving condenser water set-point on the overall performance.

9.2.2 Model Validation with Experiment Data

9.2.2.1 Validity of Models

This section discusses validity of models for cooling water generation, storage, and consumption based on comparison between experimental and simulation data for the Week 2 of Experiment 1 (the MPC policy week). The simulation data are generated with the same inputs—extracted from the building control system—that were realized for the plant as illustrated in Figure 9.15, which illustrates the data flow between several subroutines used to extract, simulate, and generate comparison plots for model validation.

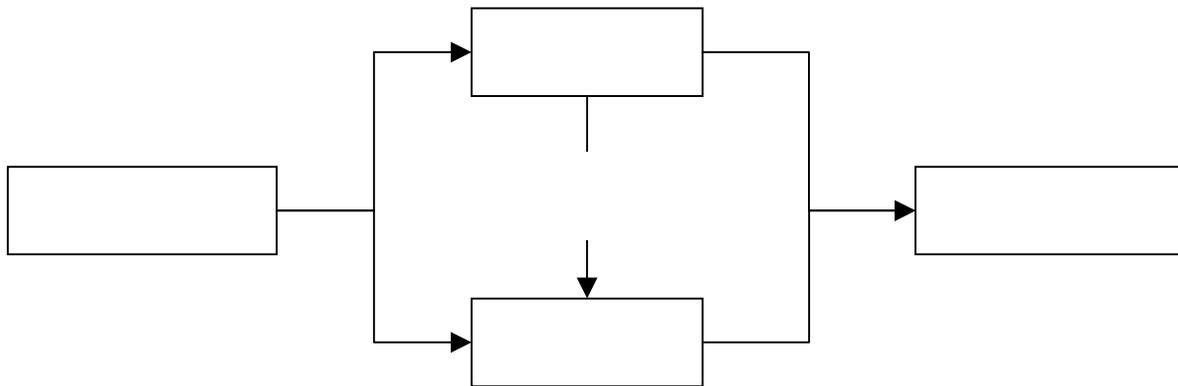


Figure 9.15. Data flow for chiller plant model validation

The implemented values of the three control variables \dot{m}_{CWS} , T_{CWS} , T_{CWS} over the entire MPC policy week are illustrated in Figure 9.5. The plotted values are the output of the closed loop systems illustrated in Figure 4.1 (Section 4.1) that have as references the optimal set-points generated by the MPC algorithm. These plots illustrate:

- The optimal set-point for the chiller water supply temperature, 1st subplot, is practically the same as the regular schedule implemented at UC Merced, approximately 3.3°C (38°F);
- The pump flow, 2nd subplot, is lower than maximum and this significantly decreased the chiller Partial Load Ratio leading to lower than targeted COP values. As mentioned previously this low pump flow values are a result of the fixed optimization interval window length (with variable starting time) and constraints on the tank capacity (a result of robust invariant set constraints driven by the campus load estimates bounds). By comparing the pump flow plot of this figure with the wet bulb temperature in Figure 9.17, one observes a non-intuitive miscorrelation between them: when the wet bulb temperature decreases, the pump flow decreases too without using efficiently the ambient conditions. This miscorrelation was corrected in the second half of the MPC policy week by increasing the number of iterations for the optimization algorithm and therefore generating an approximating solution closer to, and more representative of, the true optimal policy.

- The leaving condenser water actual values, 3rd subplot of Figure 9.60, illustrate an increase of the set-point value from 18.9°C (66°F) to 20°C (68°F). Although the optimal value generated by the MPC algorithm was higher than 20°C (68°F), the team decided to conservatively increase this value to avoid chiller surging conditions that were experienced previously with Chiller 1. The subplot also illustrates higher frequency oscillations around higher T_{CWS} (in the 2nd half of the week)—a result of mistuned PID gains for this newer values. These gains have been retuned during the 3rd week of Experiment 1. The impact of this re-tuning on the performance analysis is discussed in the next section.

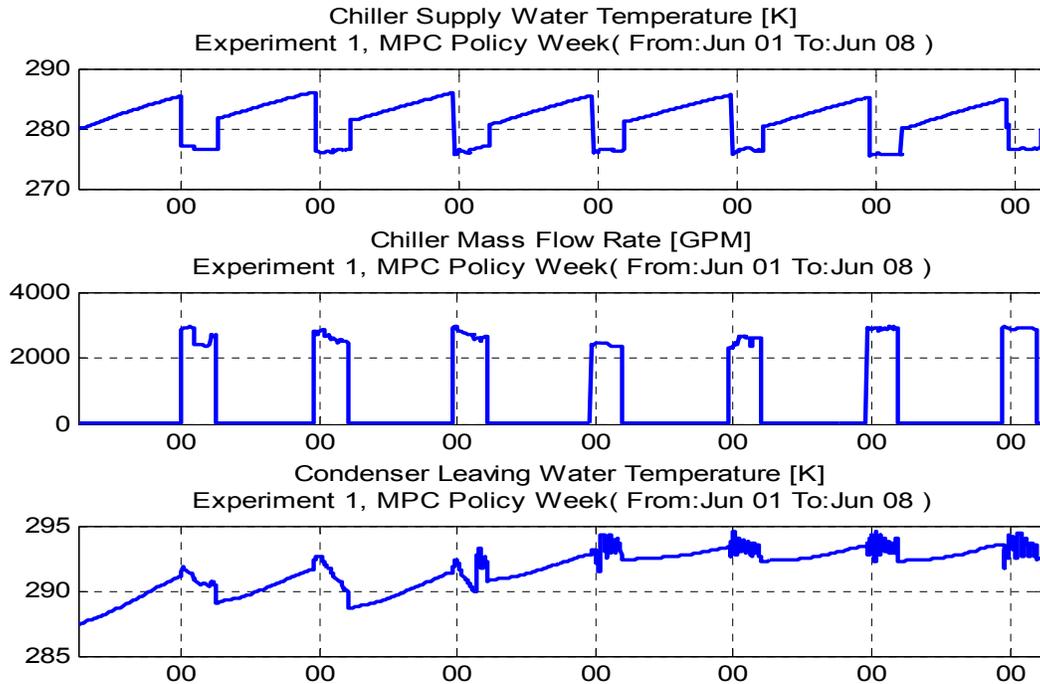


Figure 9.60. Plots of actual values of the control variables during the MPC policy implementation week of Experiment 1

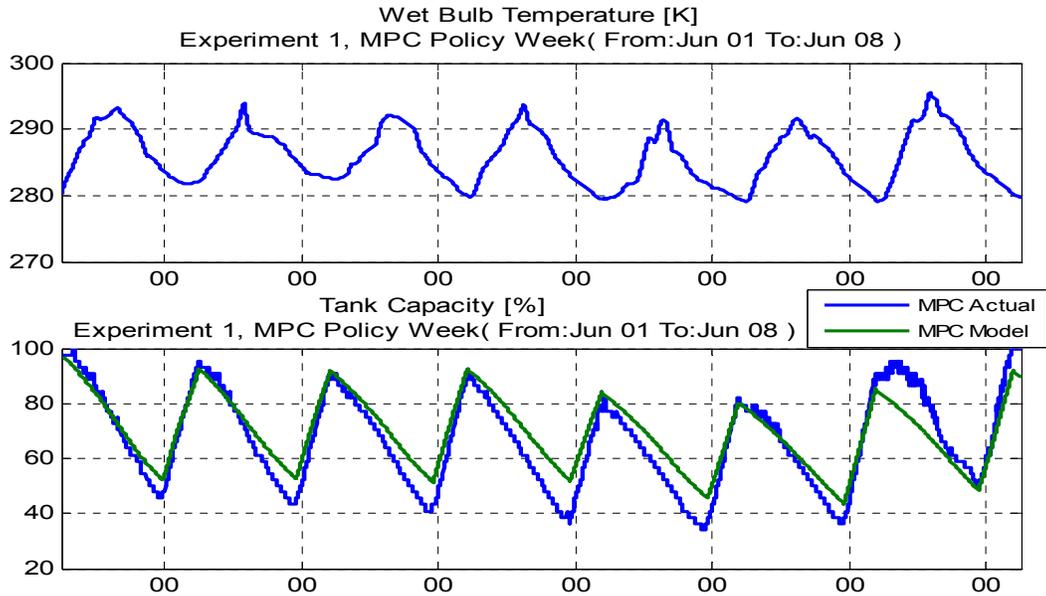


Figure 9.17. Plots of wet bulb temperature and tank capacity during the MPC policy implementation week of Experiment 1

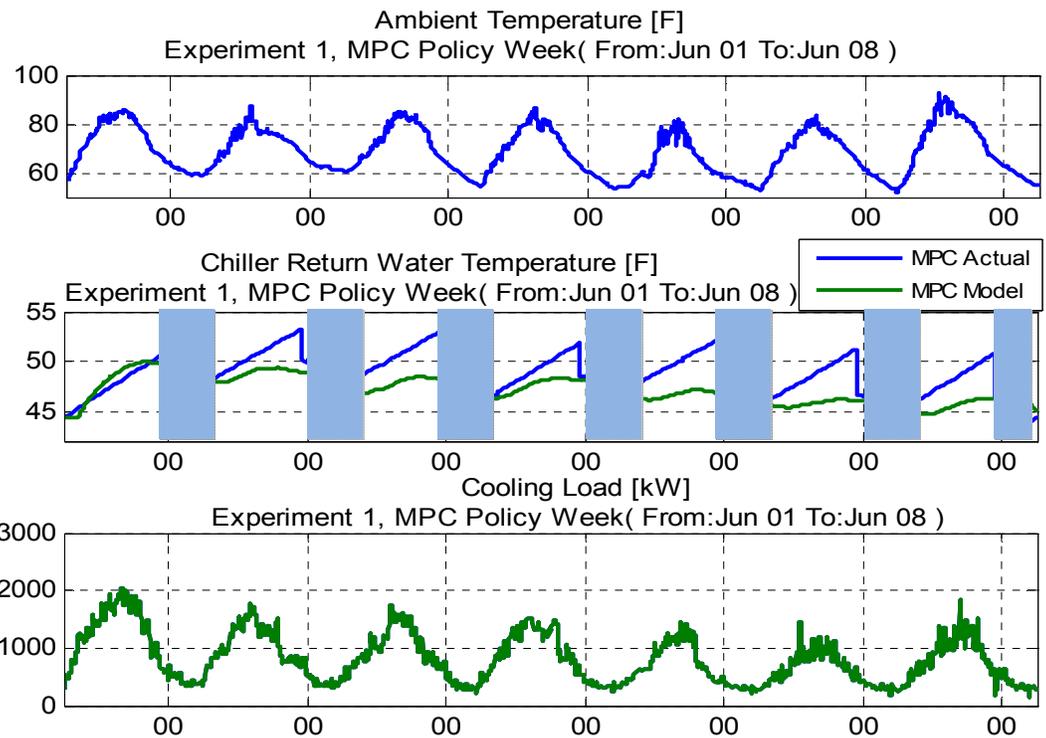


Figure 9.62. Plots of the ambient temperature, chiller water return temperature and actual cooling load during MPC policy implementation week of Experiment 1

Two of the most critical states of the model used by the optimization algorithm, and that have a

direct influence on COP, are tank capacity and chiller water return temperature plotted in Figure 9.17 and Figure 9.62, respectively. The comparisons of the model-based simulations against actual data suggest a maximum error of approximately 10% in tank capacity model and a relatively accurate average chiller water return temperature. The primary factor for the simulation-vs-measurements tank capacity mismatch is the inaccuracy in the flow rate measurements. Another factor is the reduced order model that lumps the warm and cold water stratified temperatures into only two temperature states, above and below the thermocline. The comparison illustrates higher actual charging and discharging rates than actual during the first 5 days of the MPC policy week. The error in the chiller water return temperature is relatively small during the charging window, as illustrated by the highlighted intervals of the second plot in Figure 9.62. During the last two days, the MPC policy lower bound constraints on tank capacity were increased to meet the 100% terminal constraint at the end of the implementation week.

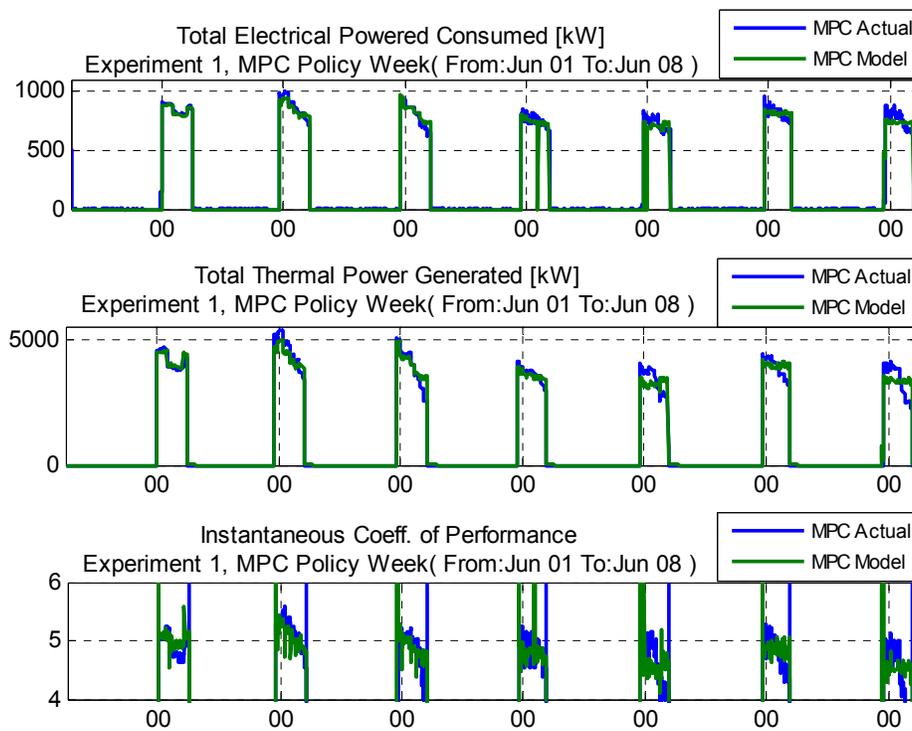


Figure 9.63 Plots of total electrical and thermal powers, and coefficient of performance during the MPC policy implementation week of Experiment 1

The coefficient of performance (COP) estimate defined previously as the thermal power (energy) produced per unit of electrical power (energy) consumed is plotted in Figure 9.63 against the actual COP determined based on time series data available from the building control system. Although Figure 9.63 illustrates instantaneous differences of up to 0.5, the average model-based estimate is a good approximation of the actual average COP. Figure 64 shows the small relative differences between total generated and consumed energies and week-long cumulative COPs. Based on these comparisons, it may be concluded that the models offer good correlation with actual plant performance, and therefore they can be used for on-line

optimization to generate optimal set-points.

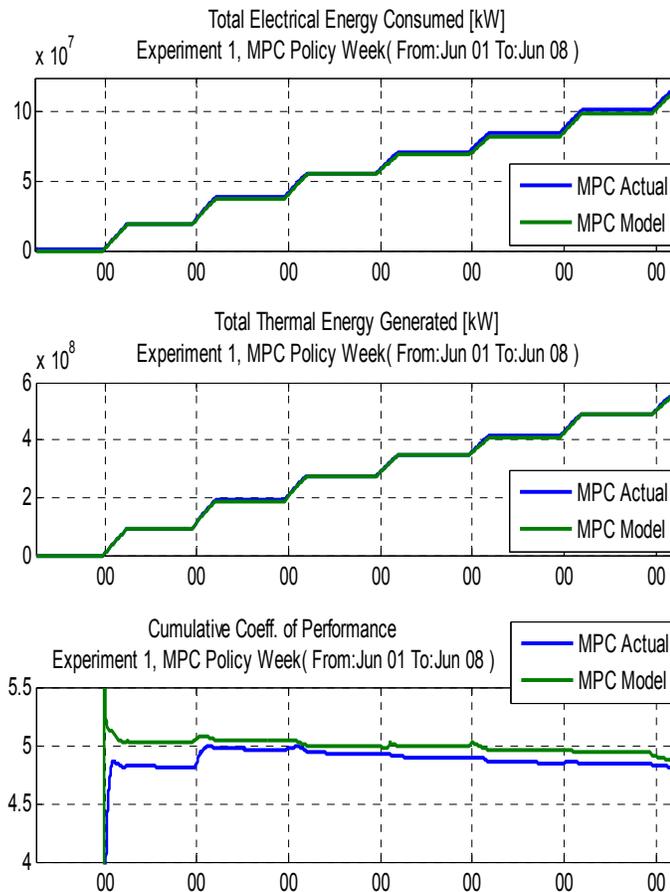


Figure 64.9. Plots of actual and simulated total electrical and thermal energies, and cumulative COP for the Experiment 1 MPC policy week

9.2.3 Performance Estimation

This section presents the time series data for all relevant signals—inputs, states, performance metrics—and discusses the differences between the MPC and baseline (regular chiller operational schedule) policies, together with their implications. This section focuses on comparison of the performance under MPC with the performance under the policy implemented during the 3rd week of Experiment 1, when the standard set-points were expected to be implemented in the building control system. This 3rd week policy is referred to as Baseline 2 or ‘adjusted baseline policy,’ due to the use of a value for the leaving condenser water set-point that was equal the MPC policy set-point, although the experiment plan was to use the original regular operation schedule.

9.2.3.1 Performance Estimation for Baseline and Experiment Weeks

The three control set-points for both MPC and Baseline 2 policies are illustrated in Figure 9.65. The three notable differences correspond to: pump mass flow rate, condenser leaving water set-point, and charging window.

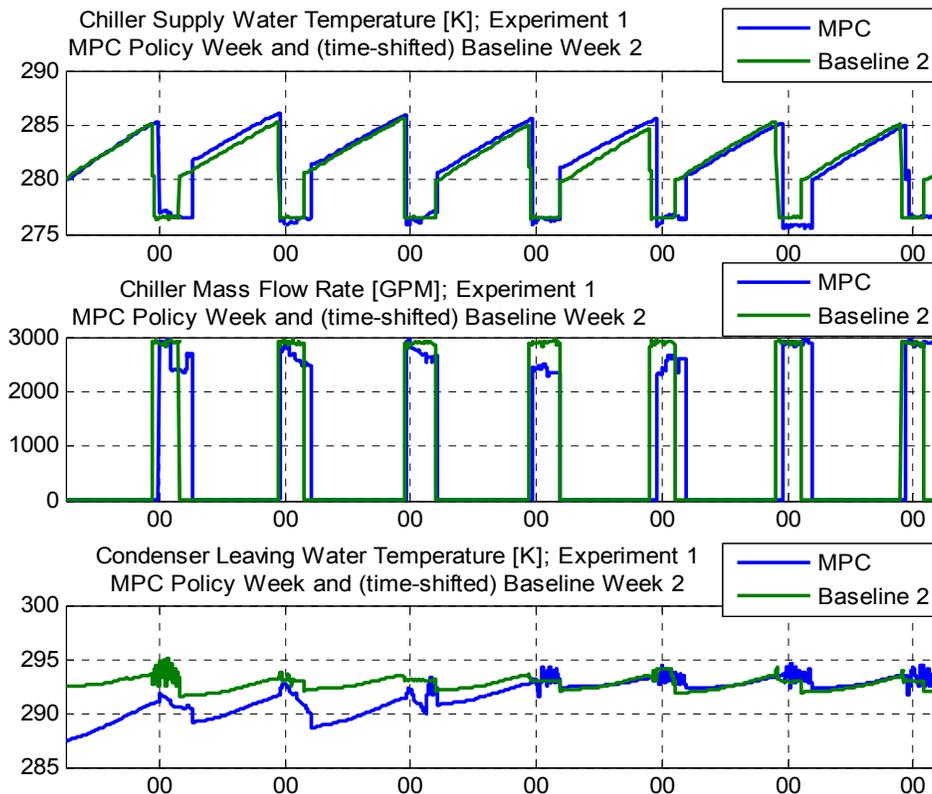


Figure 9.65. Plots of set-points values for chiller water supply temperature and flow rate for Experiment 1 MPC and Baseline week 2

The effects of these control variables, plotted in Figure 9.65, on the expected efficiency improvements mentioned in Section 4.1 are summarized below:

- Higher T_{CWS} set-point values resulted in lower cooling tower power consumption as initially predicted. This is illustrated by correlating the subplot 3 of Figure 9.65, illustrating the higher T_{CWS} set-point for the Baseline 2 week, with subplot 1 of **Error! Reference source not found.**, illustrating lower cooling tower powers for Baseline 2 week. The trend is apparent during the first two days of the respective experiment weeks. However, due to a conservative increase in the T_{CWS} during the MPC week, this improvement was not fully demonstrated during this week. The adjusted baseline schedule, used during Baseline 2 week, benefited to a larger extent from the higher T_{CWS} values.
- Lower chilled water flow rates resulted in lower chiller efficiency for similar chiller water return temperatures. During the MPC week, the optimization algorithm used a charging window with a fixed and a priori optimized length. The consequence of the fixed

charging window was a chilled water flow rate lower than the normal schedule flow rate (subplot 2 of Figure 9.65). As illustrated in Section 8.2.3.2., this lower set-point generated a low Part-Load Ratio that decreased the chiller efficiency.

- The lower-thermocline tank management strategy did not result in expected efficiency improvements. As discussed in Section 6.1, the smaller fraction of chilled water volume in tank was predicted to generate higher chiller return temperature, and in turn increase PLR and chiller efficiency. During the second parts of the two weeks, Subplot 3 of **Error! Reference source not found.** shows similar levels of chiller return temperature, for similar ambient temperature and loads, but with significantly different tank warm water volume.

[b14]

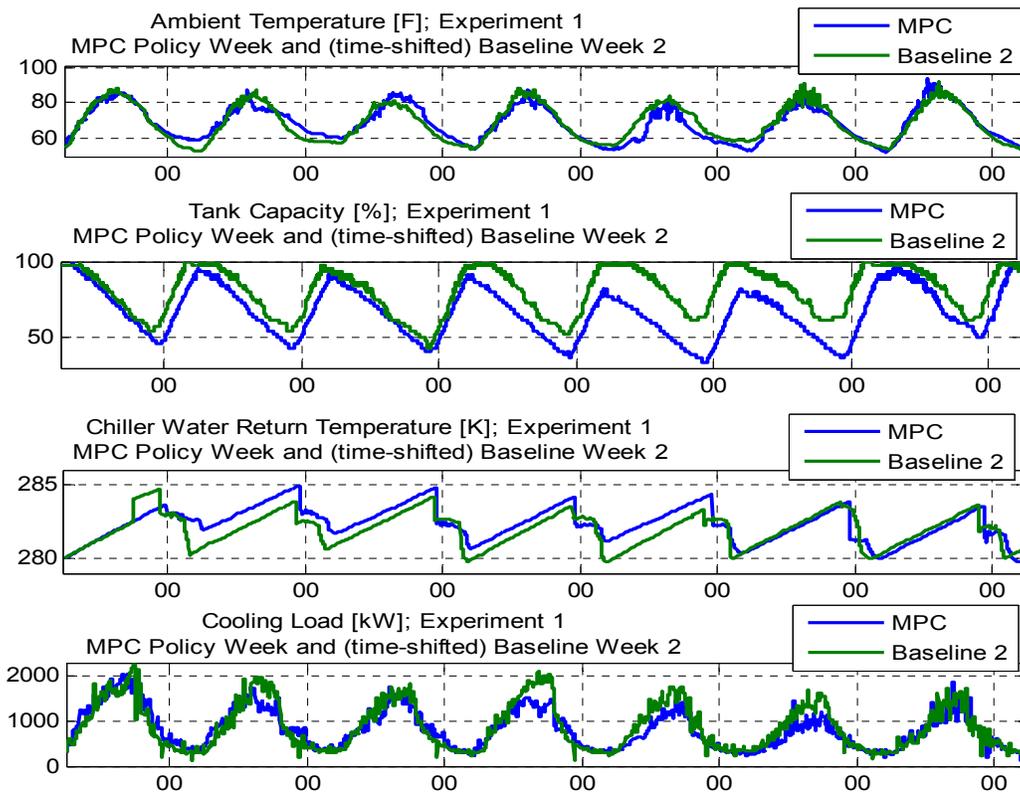


Figure 9.66. Plots of ambient temperature, tank capacity, chiller water return temperature, and cooling loads for Experiment 1 MPC and Baseline week 2

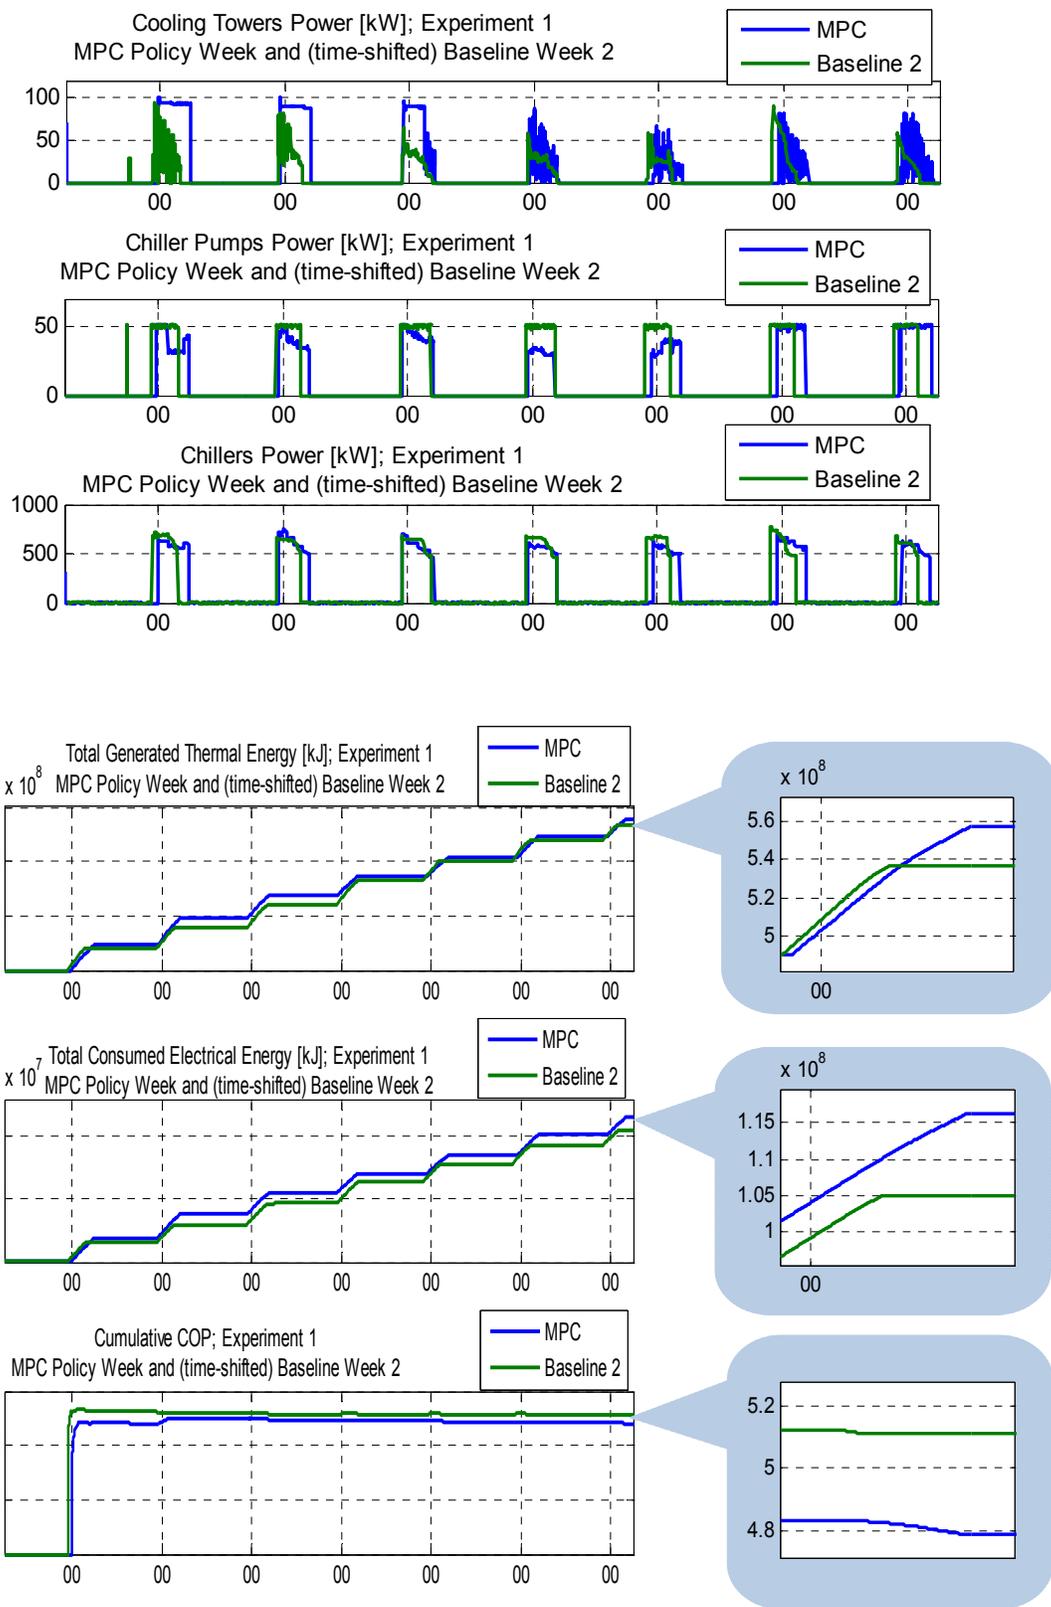


Figure 9.67. Plots of generated and consumed energy, and cumulative COP for Experiment 1 MPC policy and Baseline 2 weeks

9.2.3.2 Performance Comparison and Limits to Savings

The overall results for each of the three weeks of Experiment 1 are summarized in Table 9.5 and Figure 9.68, where the average and standard deviation values of the main variables are also presented. A correlation between these averages and the COP final values confirms qualitatively the initial chiller and tank energy performance improvements control strategy: higher T_{CHWR} and higher \dot{m}_{CHWS} , while meeting the critical constraint $PLR < 1$, result in lower average electrical energy consumption per 1kJ of generated thermal energy.

Table 9.5. Summary and comparison of main control, state, and performance variables for Experiment 1 weeks (the numbers [m s] represent average and STD, respectively)

Scenario Context		WEEK 1 (Baseline)	WEEK 2 (MPC)	WEEK 3 (Adjusted Baseline)
Weather & Campus	T_Ambient [avg,std] [F]	[77, 12]	[68, 10]	[69, 10]
	T_CHWR chiller [avg,std] [F]	[49, 3]	[48, 2]	[48, 2]
Plant Operation	T_CWS actual [avg,std] [F]	[66, 4]	[67, 2]	[68, 1]
	Mdot_CHWS actual [av,std] [GPM]	[2456, 204]	[2660, 260]	[2898, 27]
	T_CHWS actual [avg,std] [F]	[38, 1]	[38, 1]	[38, 0.1]
Energy	Energy (electrical) consumption [MJ]	175,260	116,389	105,045
	Energy (thermal) generated [MJ]	835,921	557,024	536,632
Coefficient of Performance	COP1	4.77	4.79	5.11

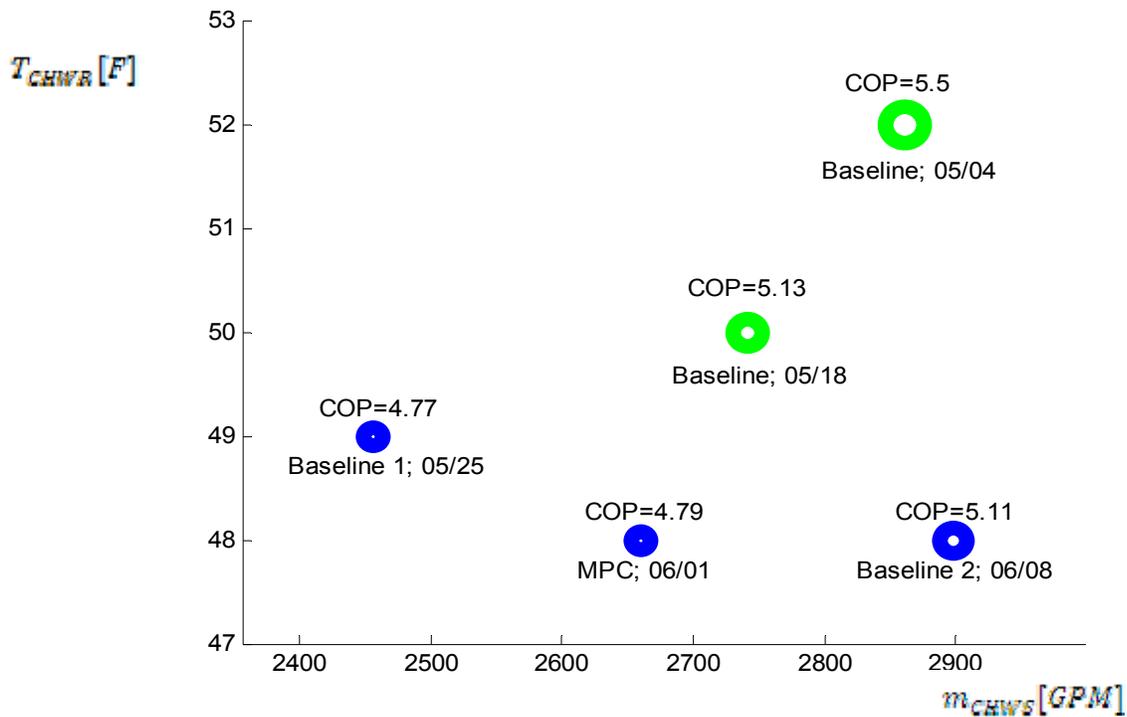


Figure 9.68. Performance comparison between baseline and MPC policy weeks. Remarks: (i) the color code is: blue for experiment weeks, and green for various other weeks; (ii) the relative size circles represents the relative size of the COP magnitudes [PH16][b17]

Based on comparisons between Experiment I COP values from Table 9.5 and Figure 9.68, it can be concluded that [b18] the performance sensitivity with respect to exogenous inputs, dependent on ambient conditions and/or campus load, is significantly higher than the performance sensitivity with respect to some of the controllable inputs. [PH19] This is illustrated by the comparison between COP values for Baseline 1 and MPC week, and likewise between Baseline 2 and Baseline for the 05/18 week. For example, for approximately the same COP values, a decrease of only 1F in T_{CHWR} , from 9.4°C (49°F) to 8.9°C (48°F), is compensated by an increase in the pump flow from 154 kg/s (2450 GPM) to 166 kg/s (2650 GPM). [PH20]

In conclusion, the MPC algorithm performance for Experiment I was approx. 6% lower than the *adjusted* baseline schedule performance, for similar average ambient temperature and humidity. To a large extent this lower performance is due to the lower values of the chilled water flow rate set-point generated by the predictive algorithm (MPC week average is 167 kg/s (2660 GPM) and Baseline 2 week average is 182 kg/s (2898 GPM)). To a smaller extent the lower performance is due to conservative values of the condenser leaving water set-point during MPC week. This new, higher, T_{CWS} value is a result of the optimization algorithm and it was used consistently during the Baseline 2 week.

This first set of experiments provided insights into the model capability to predict critical process variables and into the optimization algorithm parameters influence on these variables. These insights were used to design the chiller water supply temperature experiment and to improve the optimization algorithm, improvements demonstrated during Experiment II discussed in Section 4.3.

9.3 Appendix 3: Chilled Water Supply Temperature Experiment

9.3.1 Motivations for Experiment

The campus return temperature model (Section 2.8) produces the somewhat non-intuitive result shown in Figure 9.69: instead of always producing a higher return temperature, an increased supply temperature produces a higher return temperature at low loads, but a lower return temperature at high loads.

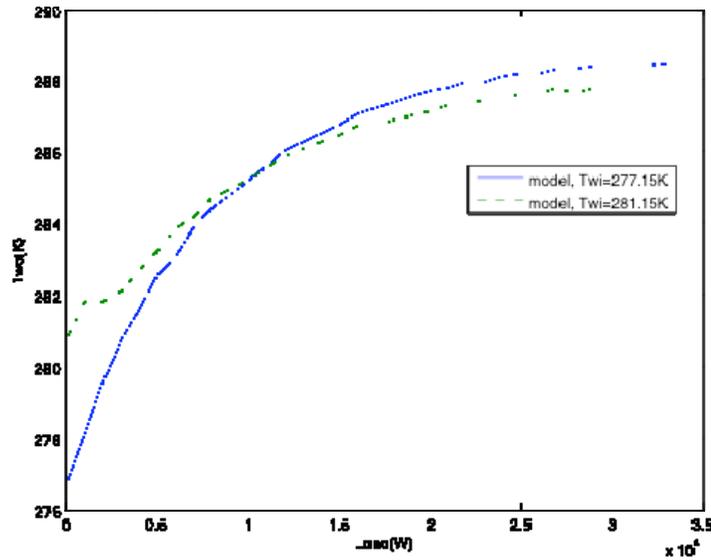


Figure 9.69: Return temperature vs load

As noted in Section 2.8, the model is solving the following two equations for T_{wo} and \dot{m}_w (all other variables are parameters or inputs or calculated explicitly from them):

$$\dot{Q}_{load} = \dot{m}_w c_{p_w} (T_{wo} - T_{wi}) \quad (9.2)$$

$$\dot{Q}_{load} = UA\Delta T_m = \zeta \left(\alpha_1 \dot{m}_w^{0.85} + \alpha_2 \dot{m}_a^{0.95} \right)^{-1} \left(\frac{T_{ao} - T_{ai} - T_{wi} + T_{wo}}{\log \left(\frac{T_{ao} - T_{wi}}{T_{ai} - T_{wo}} \right)} \right) \quad (9.3)$$

Equation 9.2 suggests that the return temperature would always increase with an increase in the supply temperature, and for a given water flow rate and load, this is what happens with this equation. However, for that same given flow rate and load, Equation 9.3 calculates a lower return temperature, for the following two reasons: (a) for given mass and air flow rates, the UA is given, thus the LMTD is a given; (b) to maintain a given LMTD when the inlet water

temperature is increased, the outlet water temperature must be decreased, as illustrated in **Error! Reference source not found.**

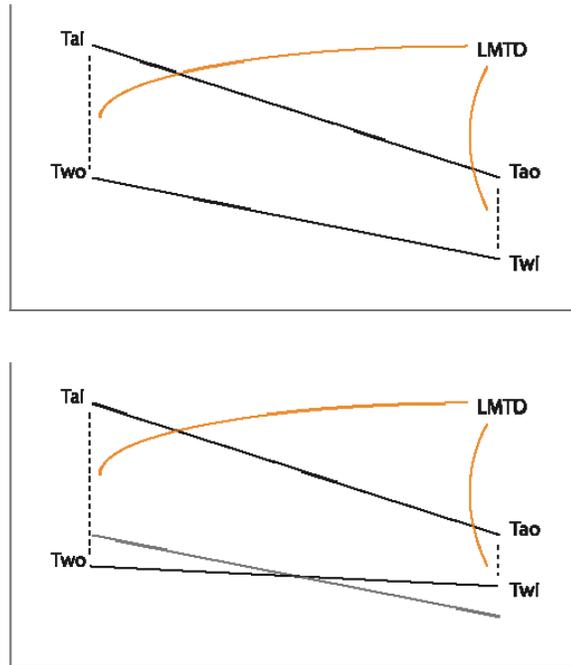


Figure 9.70: Maintaining a Log Mean Temperature Difference

So the effect of increasing the supply water temperature is to push the solution towards higher water flow rates, and the new intersection between the two equations may result in either a higher or lower return water temperature. At high loads (**Error! Reference source not found.**), the effect is to decrease the return temperature, whereas at low loads (**Error! Reference source not found.**), it increases the return temperature. A possible rationale for this: at lower loads the slope of the UA vs \dot{m}_w curve is steeper than it is at higher loads, and thus the slope of the second equation is steeper at lower loads than at higher loads. Note that this result is highly sensitive to the values of the calibrated α_1 and α_2 values - at some combinations of these values the effect of increasing T_{wi} is always to increase (or decrease) the solution value for T_{wo} .

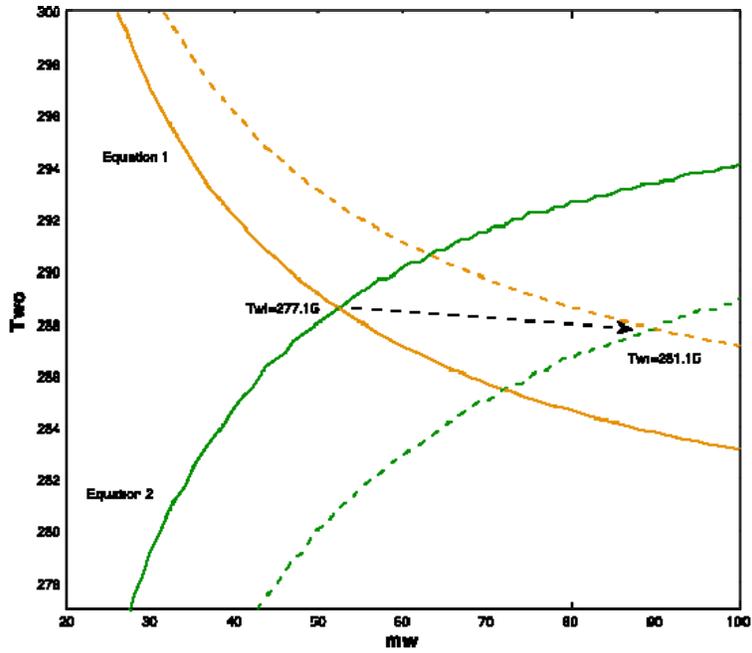


Figure 9.71: Effect of increasing T_{wi} at Load = 2.5MW

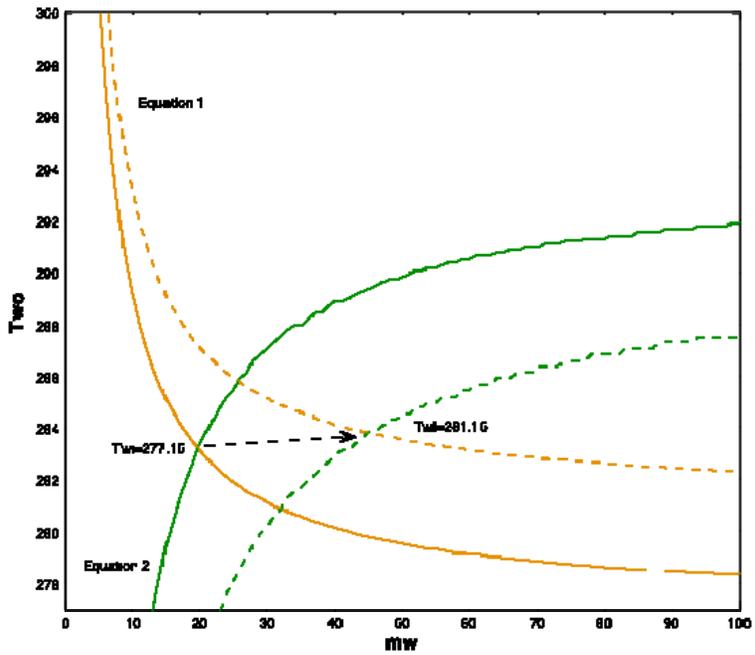


Figure 9.72: Effect of increasing T_{wi} at Load = 0.5MW

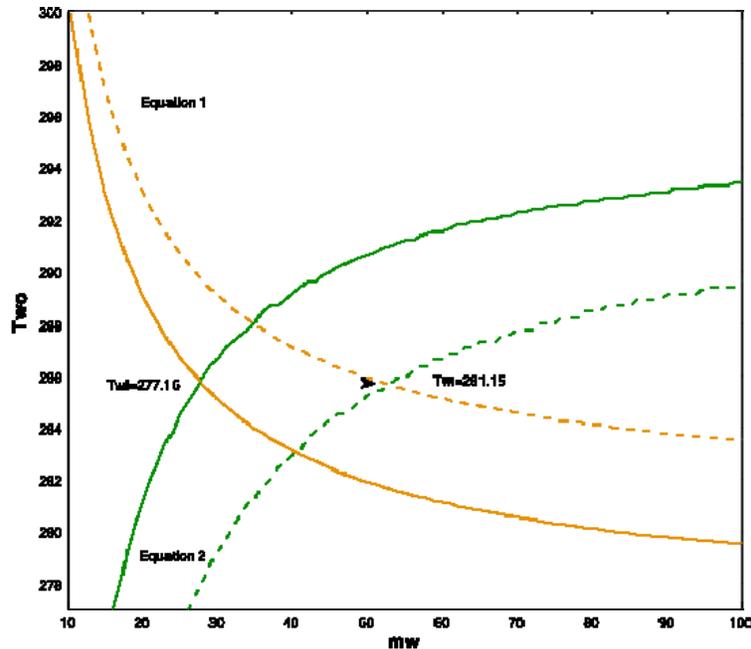


Figure 9.73: Effect of increasing T_{wi} at Load = 1.0MW

9.3.2 Design of Experiment

The objective of the experiment was to evaluate the effect of a small increase in the chilled water supply temperature on the return temperature and flow rate from the campus. Because the standard operation of the chillers maintains a constant supply water temperature, there were no empirical data available to verify or calibrate the campus return temperature model outputs when supply water temperature was higher.

Before the experiment, the desired outcomes were noted as follows:

- Validate and further calibrate the campus return temperature model.
- Contribute to a more complete understanding of how chilled water supply temperature affects energy consumption of chilled water plant, pumps and AHU fans.
- Work towards the longer-term objective of identifying potential scenarios where it may be advantageous to increase the chilled water supply temperature.

The experimental procedure was as follows:

- Day 1 (9/5/2009): Fully charge tank with chilled water such that chilled water stored in the tank is 42°F. Note buoyancy of the water will facilitate mixing of water up to the level containing 42°F water.
- Day 2 (9/6/2009): Monitor critical loads. Load monitoring shall consist of monitoring indicators such as supply air temperature set-point tracking and combined cooling coil valve, pump, and fan saturation. Reduce chilled water supply temperature appropriately

if cooling loads unexpected go unmet.

- Day 3 (9/7/2009): Continue monitoring critical loads. Collect data for cooling coils subject to reduced loads.
- Day 4-6 (9/8-10/2009): Continue to monitor critical loads. Collect three days of data for cooling coils subject to typical summer loads.

9.3.3 Execution of Experiment

The experiment was carried out as planned. Critical loads were monitored throughout the experiment to ensure that no problems arose. The chilled water tank reached a stable and well-mixed temperature of 5.6°C (42°F) faster than expected. Data collection occurred without any problems.

9.3.4 Experimental Results and Analysis

9.3.4.1 Results: Campus Level

Figure 9.74 shows that there was in fact a higher supply temperature to the campus during the experiment week. (The spread in the previous week is interesting. Various possible causes were identified in discussion with the operators, but no definite conclusions were reached.) The return temperatures to the central plant were higher than expected (

Figure 9.75). The coil model was predicting that the return temperatures during the experiment week would be very similar to the previous week, with just a very slight rotation so that the return temps would be very slightly higher at low loads and slightly (about 1F) lower at high loads (

Figure 9.76).

The graph of flow rate versus load (Figure 9.77) shows something quite unexpected - the flow rates at low loads were actually lower than they were during the previous week. Why this is so unexpected is discussed in more detail in the Analysis subsection below, but the basic reason for surprise is that for a given small load, the water side delta T was higher during the experiment week than it was during the previous week, even though flow rate (and thus, one would expect, the UA) was lower during the experiment week. The model principles suggest that an increase in supply water temperature should always result in an increase in water flow rate, but this was not happening during the experiment week.

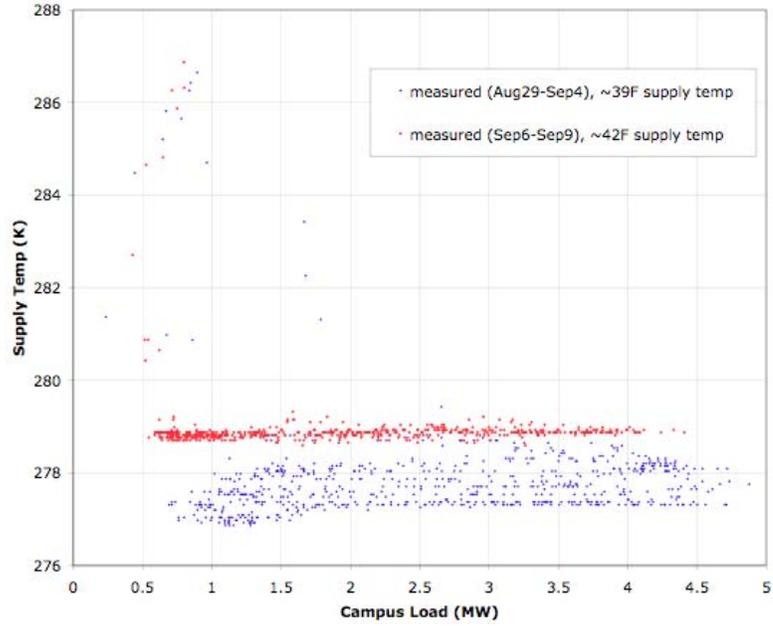


Figure 9.74. Campus supply temperature (pre-experiment in blue, experiment week in red)

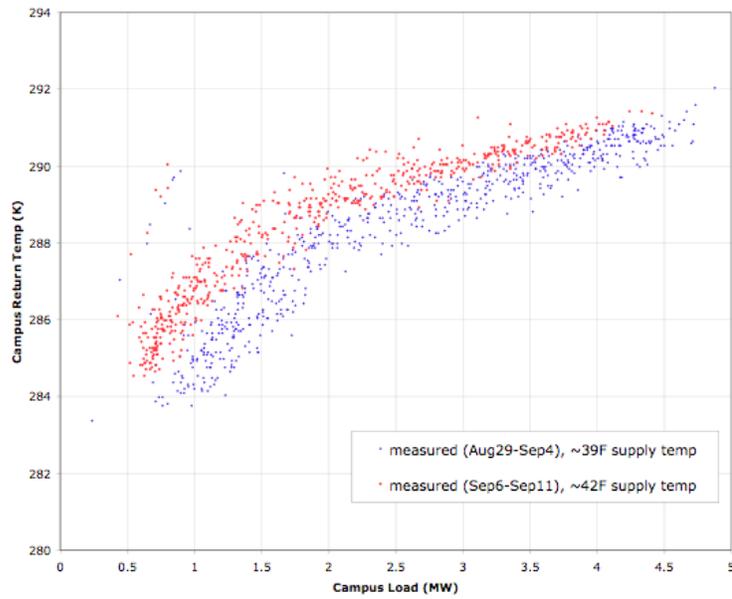


Figure 9.75. Return chilled water temperatures to central plant (pre-experiment in blue, experiment week in red)

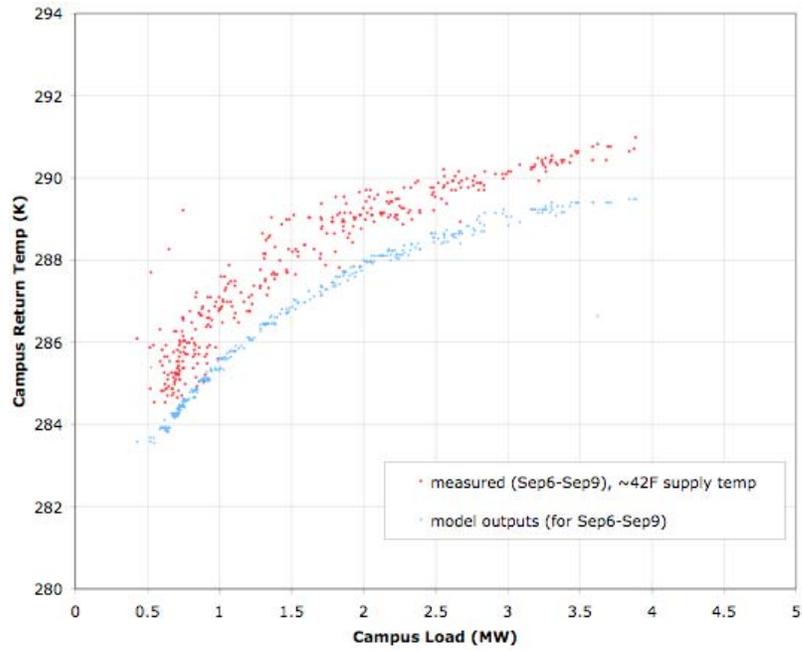


Figure 9.76. Return temps to central plant (model prediction in light blue, experiment week in red)

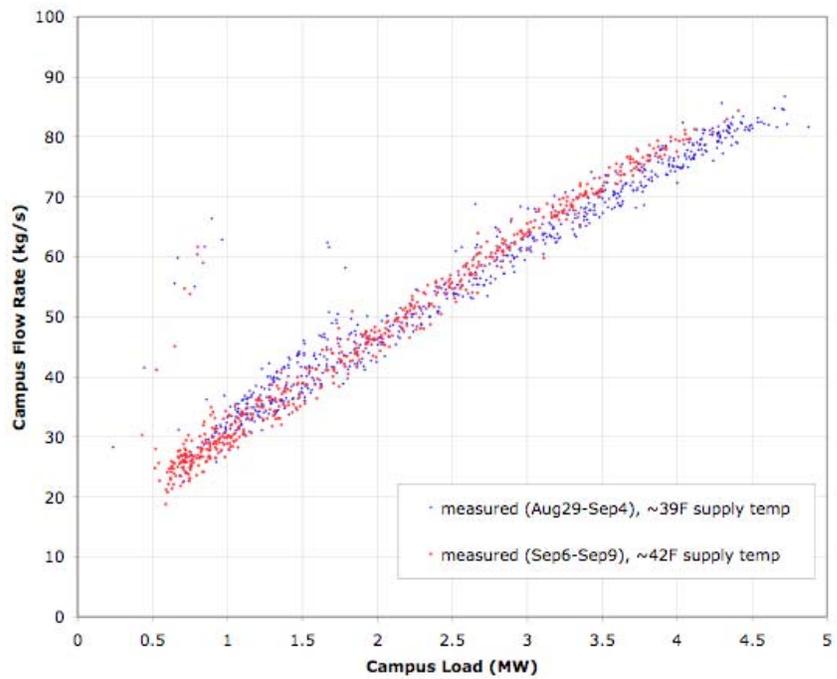


Figure 9.77. Total campus flow rates (pre-experiment in blue, experiment week in red)

Because the results are so unexpected, the details of each building's response is of interest, to see if this effect is found throughout or if there are some buildings that might be showing particularly unusual results.

9.3.4.2 Results: Individual buildings

In the results for individual buildings, the spreads in the graphs tend to be bigger, but trends are usually still visible. The chilled water return temperature model was originally developed by considering the Classroom and Office Building in detail, because it has just one air-handling unit. As shown in Figure 9.78, this building's behavior over the two weeks is similar to what we was expected: the return temperatures during the experiment week were nearly equal to the previous weeks, and on average possibly slightly higher at low loads and slightly lower at high loads.

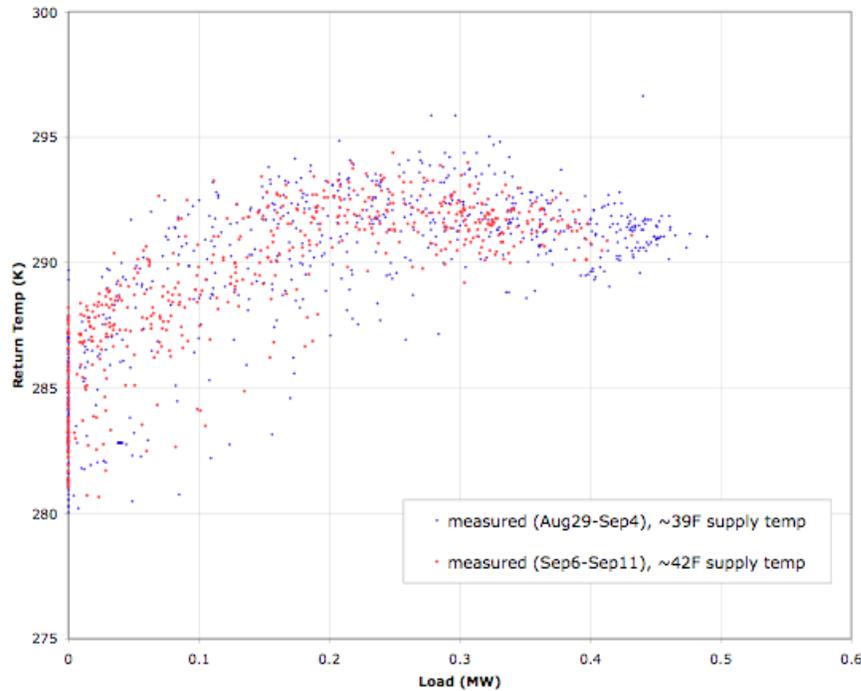


Figure 9.78. Classroom and Office Building: Return temps (pre-experiment in blue, experiment week in red)

The corresponding flow rates were as expected (Figure 9.79), with the experiment week flow rates always being slightly higher than the pre-experiment week, and showing a slightly steeper slope.

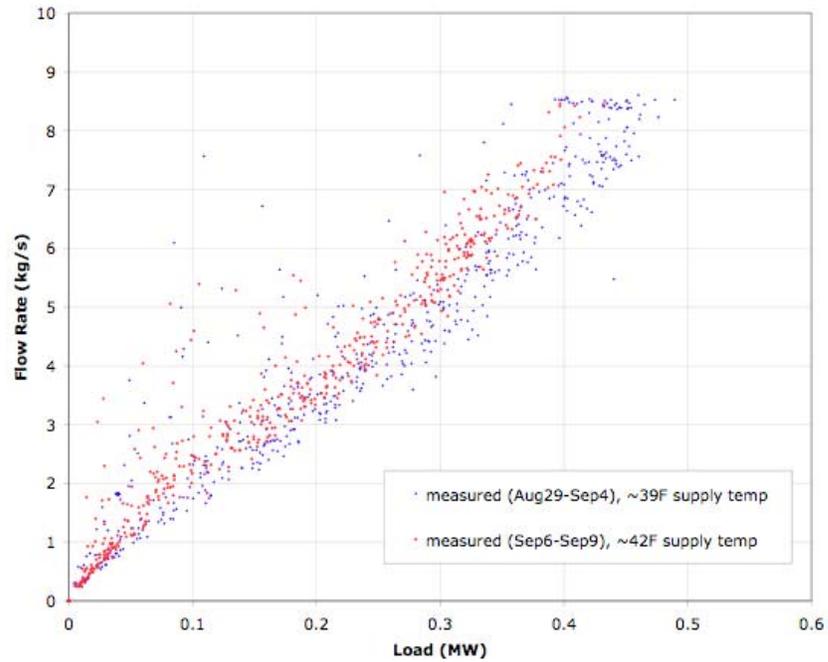


Figure 9.79. Classroom and Office Building: chilled water flow rates

The behavior of the Library Building (Figure 9.80) was also as expected, as was the Recreation Building (Figure 9.81).

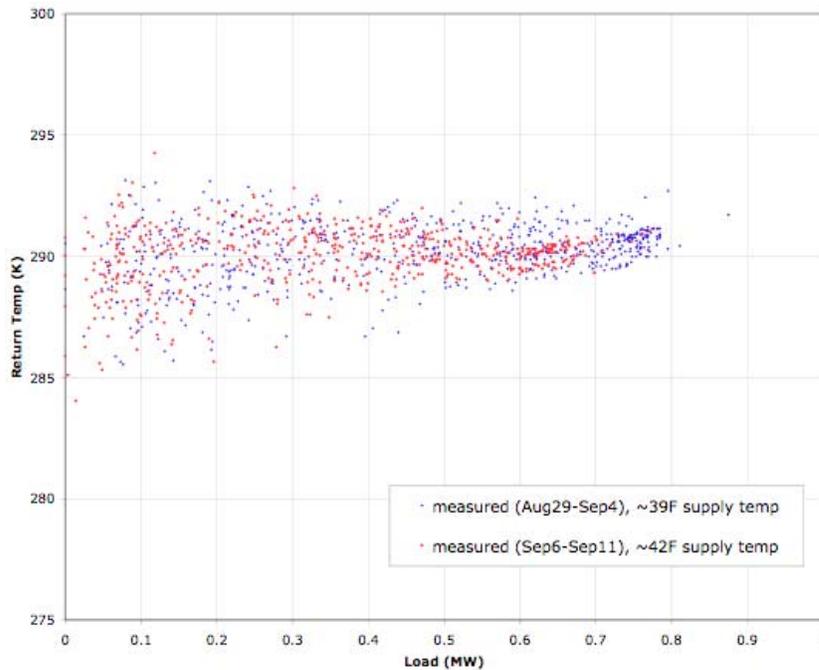


Figure 9.80. Library Building: return chilled water temperatures

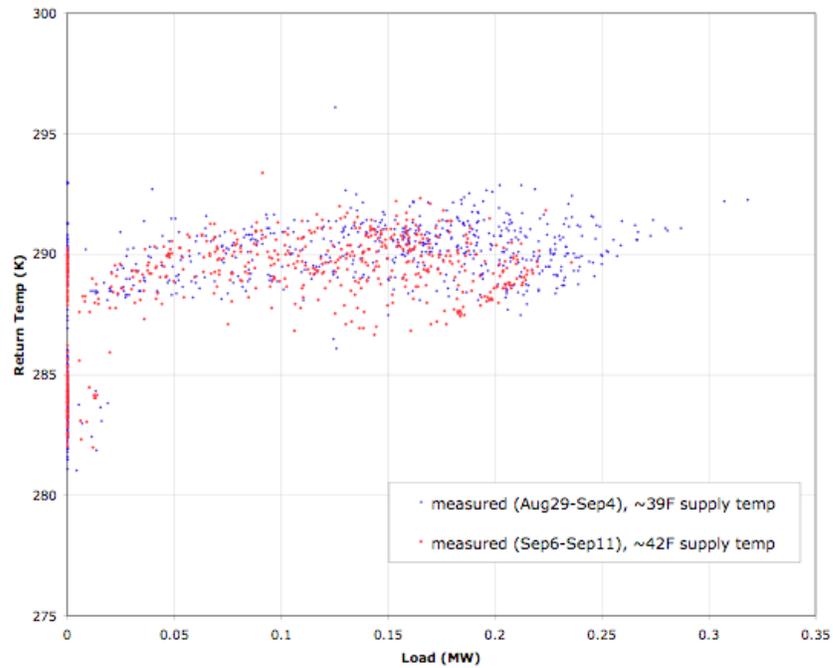


Figure 9.81. Recreation Building: return chilled water temperatures

The behavior of Sierra Terraces (Figure 9.82) was generally as expected, except at high loads. This building's data also shows recirculation at high loads, which might explain the results below, and which suggests that the bridge flow controls are set up poorly for this building and that it is not meeting set-points at high loads. It is likely not, however, affecting the general results of this experiment very much because this load is relatively small.

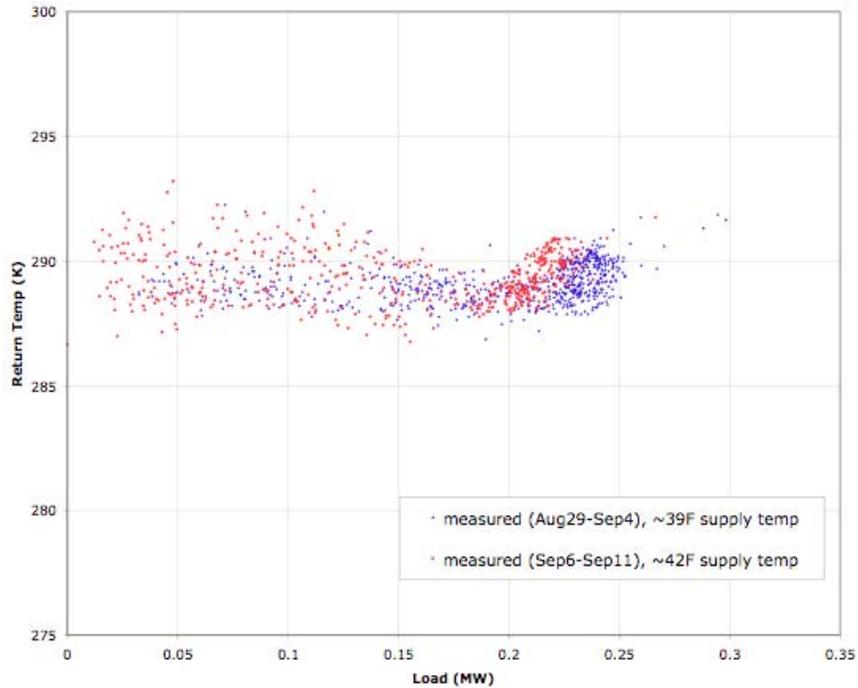


Figure 9.82. Sierra Terraces: return temperatures

The dormitory buildings are controlled differently than the rest of campus. Supply chilled water is recirculated at the bridge to maintain a desired return temperature. Although there is some minor variation from building to building, the graphs shown here for the Kings Dorm building (Figure 9.83 and Figure 9.84) are typical.

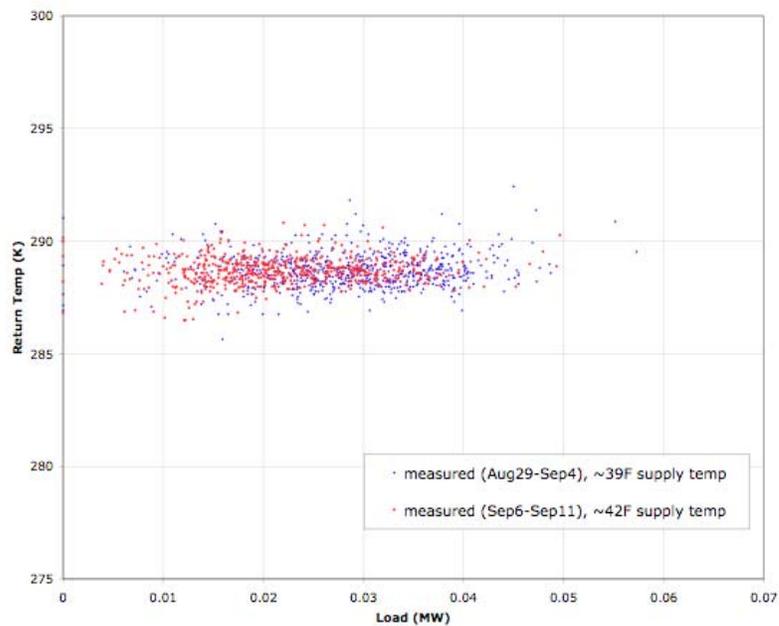


Figure 9.83. Kings Dorm: return temperatures

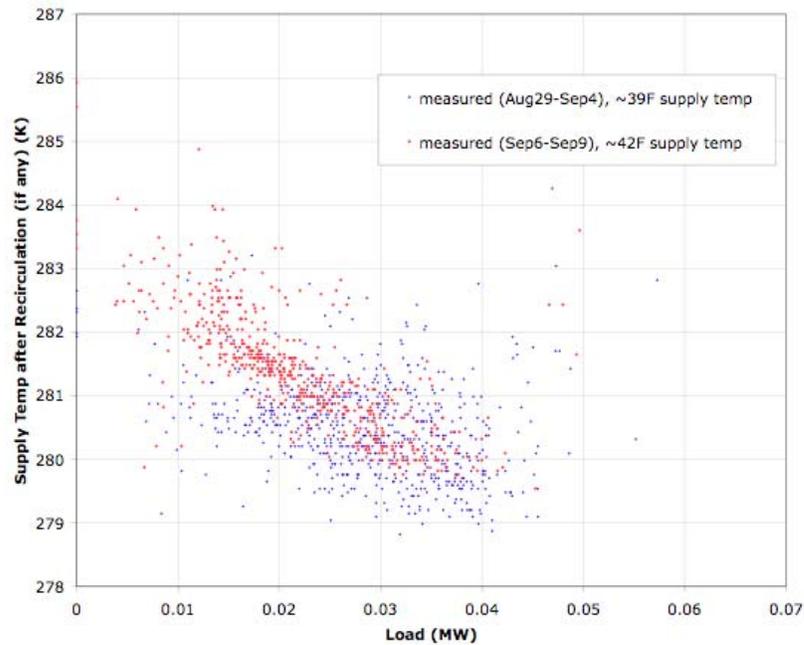


Figure 9.84. Kings Dorm: Supply temperatures after recirculation

So most of the buildings were behaving as expected and/or were producing very similar return temperatures during the experiment week as during the week prior. But there were a few that were not, including the Science and Engineering Building, which represents 25-30% of the campus load. Figure 9.85 shows the return temperatures for this building. This might be the main source of difference in the campus return temperatures between the two weeks. Various possible causes of this unexpected behavior have been discussed during the project, but further is required to fully explain it.

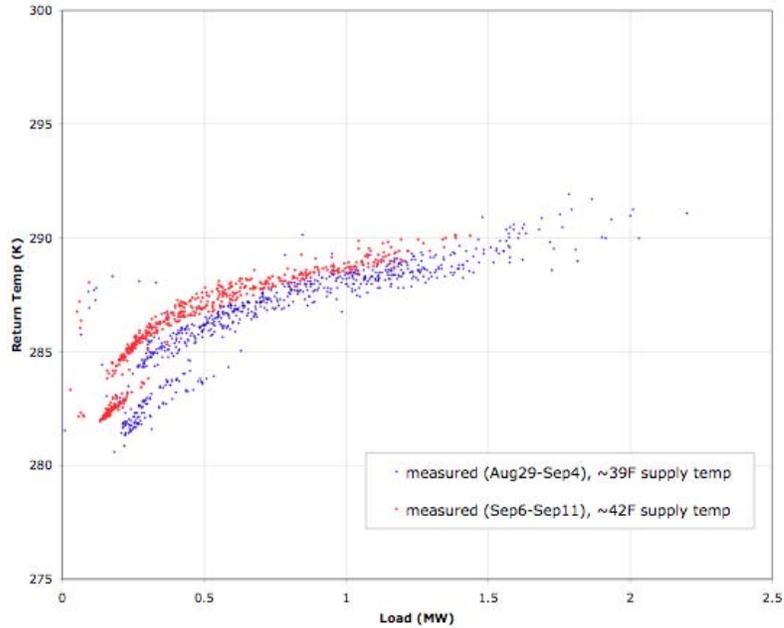


Figure 9.85. Science and Engineering Building: Return temperatures

The flow rates (Figure 9.86) at the Science and Engineering Building did not, however, seem to drop lower during the experiment week than during the previous week (unless there is something unusual happening with that second mode at low loads). On average, the loads were lower all across campus during the experiment week, but the decrease in loads was more pronounced at this building - this may have caused some of the measured effects at the aggregate level, and/or it may be pointing to something strange happening in the bridge controls for this building.

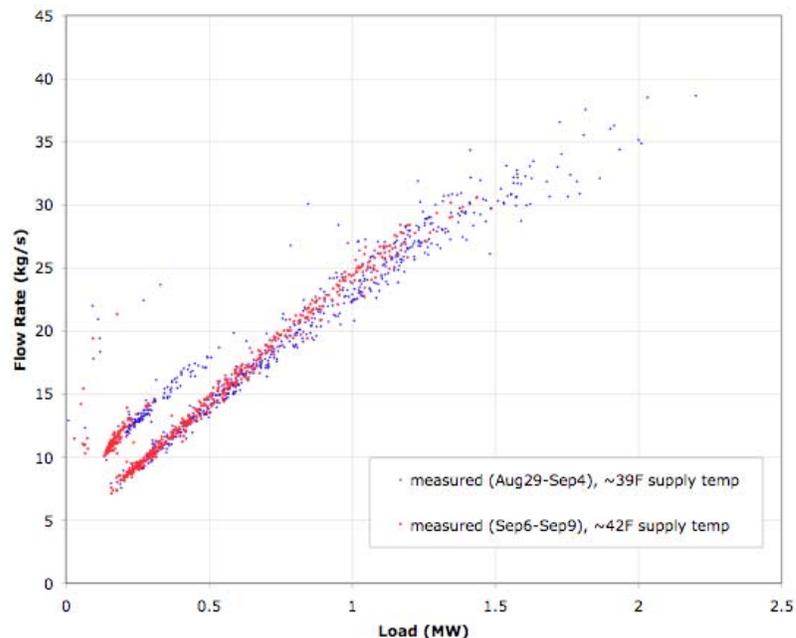


Figure 9.86. Science and Engineering Building: Flow rates

An upwards offset in the chilled water return temperature was also observed in three smaller buildings, each of which is a bit unusual in some way. The Dining Building (Figure 9.87), at about 8% of the campus load, is likely also having a significant effect on the aggregate return temperature difference. The other two are the Commons Building (Figure 9.88), at about 2% of the campus load, and the Facilities Building (Figure 9.89), at 2-3% of the campus load.

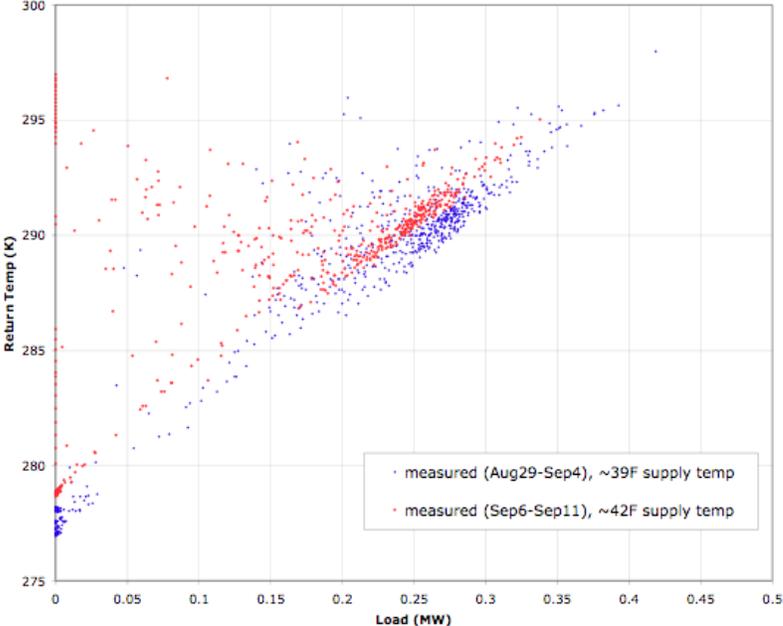


Figure 9.87. Dining Building: return temperatures

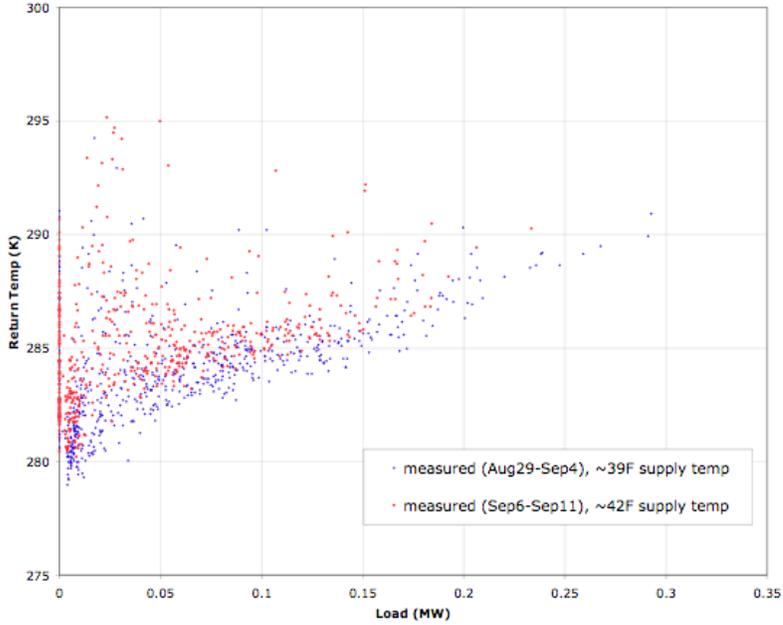


Figure 9.88. Commons Building: return temperatures

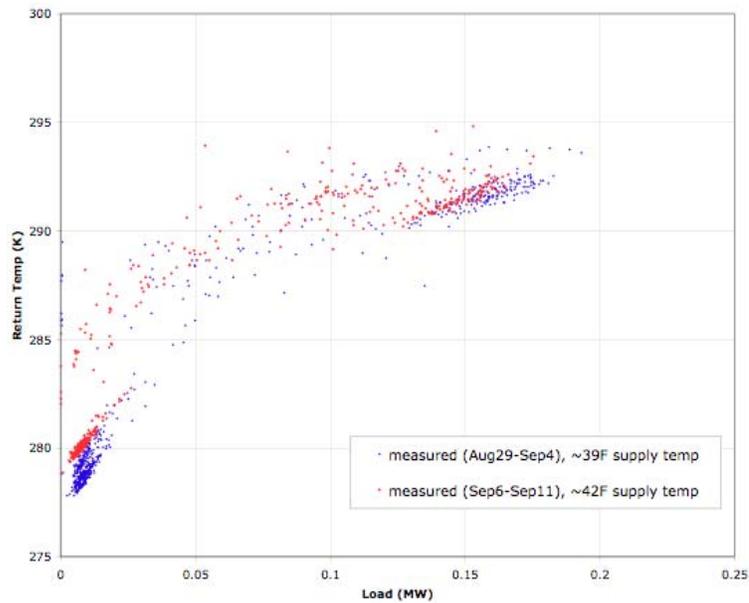


Figure 9.89. Facilities Building: return temperatures

In general, most of the buildings behaved as expected, but the Science and Engineering Building and the Dining Building are showing an upward offset in the return temperature, and their loads are big enough to possibly explain the offset for the campus loop as a whole. Further analysis is required to explain the causes of this behavior. However, for the purposes of the second MPC experiment, the campus return temperature model had to be modified to fit the empirical data. This modification process is described below.

9.3.4.3 Analysis: Model modifications to fit data before second MPC experiment

Figure 9.90 shows the measured data for the week before the chilled water experiment and from the week of the experiment. This is the data that the coil model was required to fit.

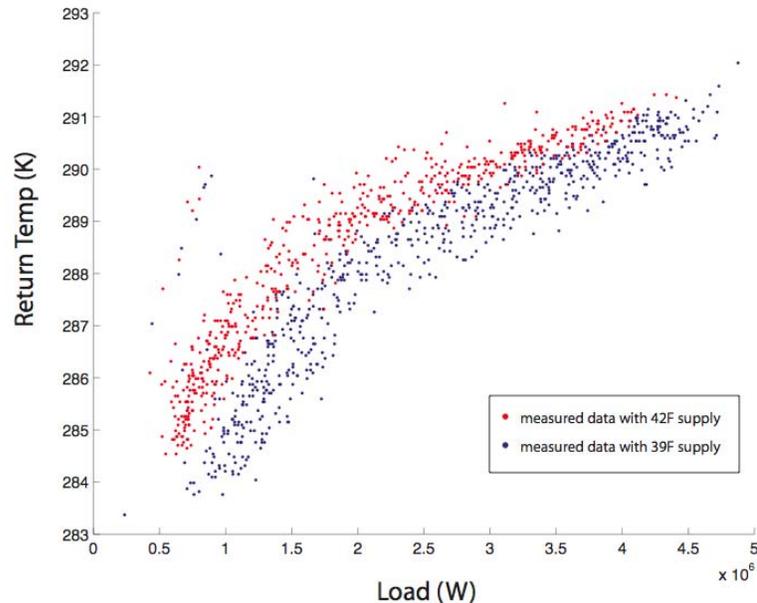


Figure 9.90. Measured Data from Experiment

The first attempt was to stretch the model to see if it could be made to fit this new data. If there was a significant weather difference during the experiment week that was not captured by the model, then the model might be able to capture the data by changing the γ parameter for the outside air fraction. Otherwise, the α_1 , α_2 and β parameters could be modified to try to capture the new data.

Figure 9.91 shows the outdoor temperature vs load for the experiment week (red) and the week before (blue). The γ (average outdoor air ratio the AHUs) parameter generally captures the spread in the return temp vs load graph (e.g. Figure 9.16), but it was conjectured that if there was a difference in weather between the two weeks, a higher γ value may also help to explain some of the increase in return water temperature during the experiment week. This seems a reasonable possibility, since the buildings that showed the increase in return temperature (e.g. the SciEng building) were ones which generally have more outdoor air requirements. By using a linear fit between the load and the outdoor temperature, and ignoring the spread, we can see if any trend change could have the speculated impact on the model.

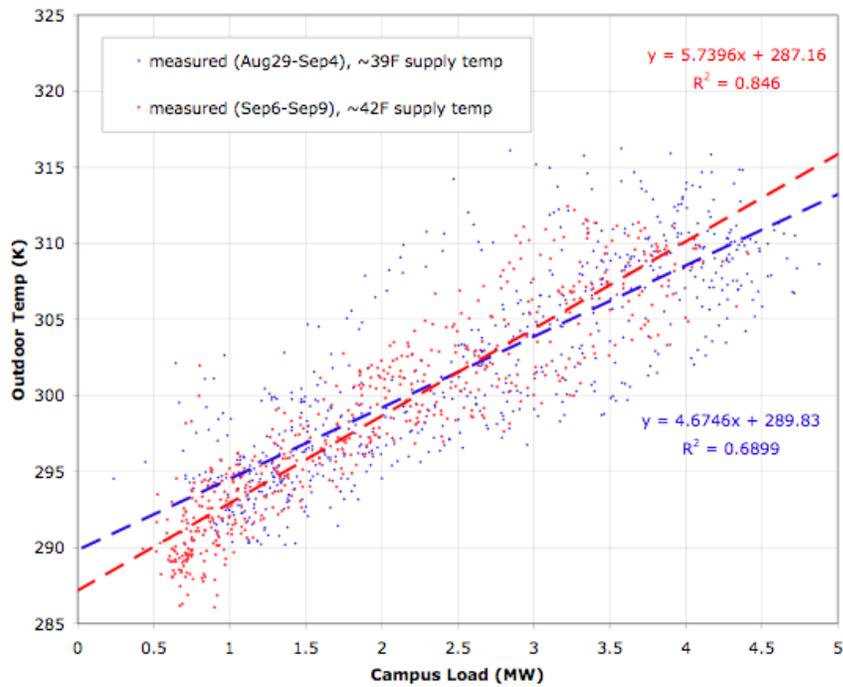


Figure 9.91. Measured outdoor air temperature vs campus load

Using the linearized T_{amb} vs Load relationship for the two different weeks, Figure 9.92 shows the effects of changing γ . In this graph and the ones that follow, like colors show the model results with the same set of parameters, with solid lines showing the results with the lower supply temperature and dashed lines showing the results with higher supply temperature. The figure shows that changing γ can change the shape of the curve between the two weeks, but cannot significantly affect the relationship between the results with the two different supply temperatures.

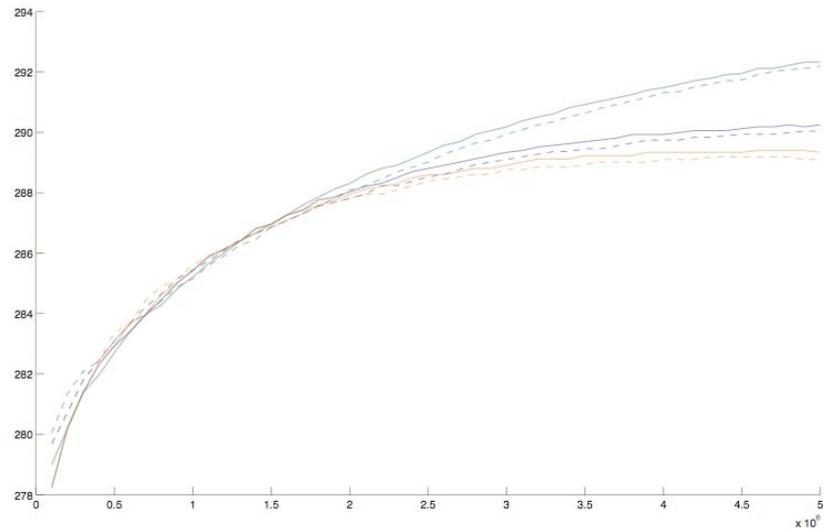


Figure 9.92. γ sensitivity

Since γ itself was unable to account for the bulk of the difference, the next step was to see if it could be accounted for by changes to α_1 and α_2 (involved in the relationship between UA and the water flow rate and air flow rate respectively) and/or changes to β (which scales the empirical relationship between UA and Load that was added earlier to account for the dropoff in UA at low loads presumably due to less smooth control at the bridges and AHUs at these loads). Figure 9.93 shows the impact of changing β - it can have some difference at low loads, but not at medium or high loads, and any changes to it greatly affect the shape of the resulting curves. So β alone cannot account for the change, but it may have some impact.

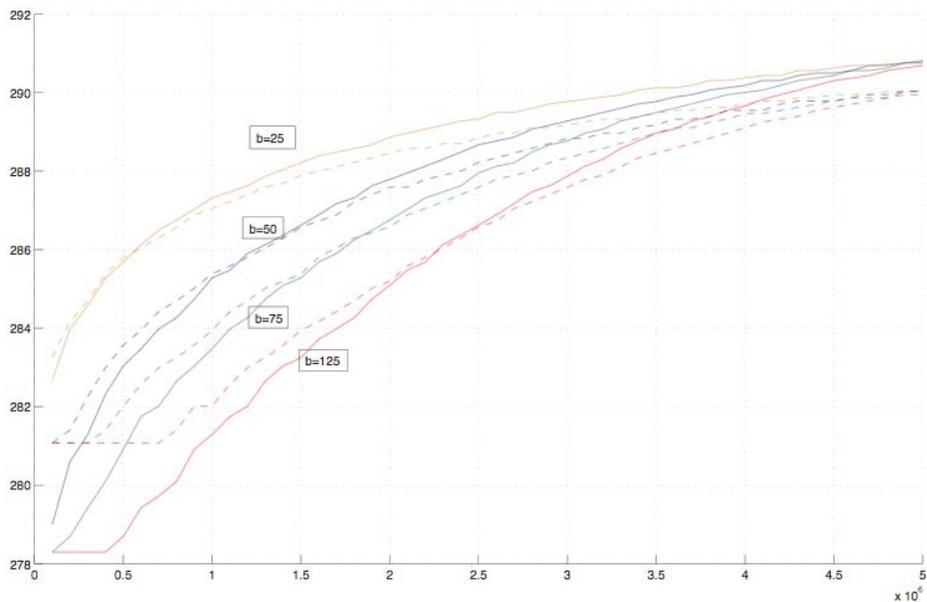


Figure 9.93. β sensitivity

Increasing α_1 and decreasing α_2 does, however, go some way to explaining the results. This changes the relative sensitivity of the UA to changes in the water and air mass flow rates: increasing α_1 and decreasing α_2 makes smaller increases in the water mass flow rate produce the same increase in UA, which allows the water side delta T and the air-water LMTD to be higher in the converged solution. Figure 9.94 shows the model stretched to its limits, with α_1 much larger than α_2 (this much difference makes the impact of the water flow rate very much outweigh the impact of the air flow rate), γ increased, and β made as high as possible without disfiguring the curve too much.

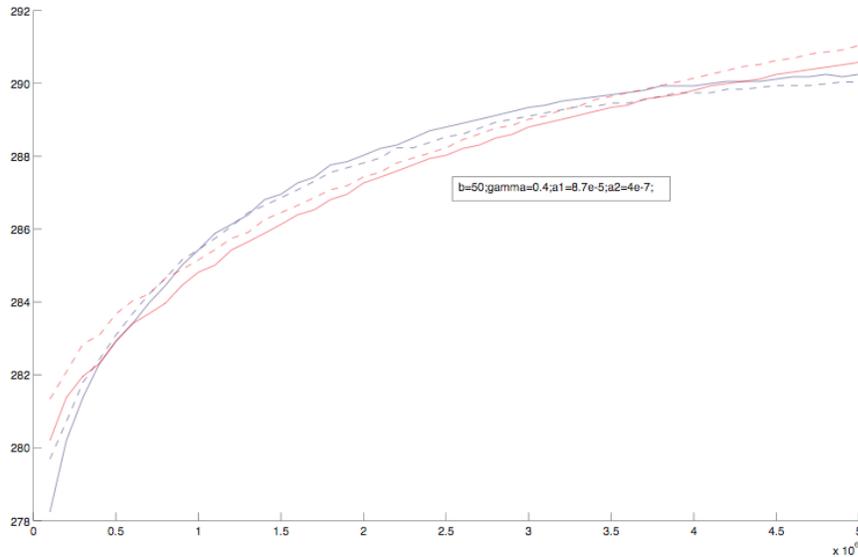


Figure 9.94. Fully stretched model (red) vs the model as configured before experiment (blue)

The fully stretched model looks like it might work, but when compared with the experimental data, it still does not produce enough of an offset between the experimental week and the week before. This result is interesting. In looking at the difference between the two weeks in red in Figure 9.94, there is a ~ 1 - 1.5 K (~ 1.8 - 2.7°F) offset at low loads, and the offset generally decreases as the load increases. This ~ 1 - 1.5 K (~ 1.8 - 2.7°F) offset is essentially the same as the change in the supply temperature. At its limit, the model is only able to produce a change in return temperature less than or equal to the change in the supply temperature. This makes sense from first principles, as an increase in water side delta T (i.e. $T_{wo} - T_{wi}$) and an unchanged load would imply a decrease in water flow rate, but if the air side remains unchanged this increase in water side delta T would require an increase in UA, which would require an increase in flow rate (according to the relationship described above). So as long as the air side values are equal between the two cases, the model must always converge to a solution where any increase in T_{wi} coincides with a smaller increase (or potentially a decrease) in T_{wo} . Or phrased differently, the water side delta T ($T_{wo} - T_{wi}$) must always decrease given any increase in supply temperature. And the relative size of this decrease is dependent on the flow rates, and thus on the load - at higher flow loads, the UA is less sensitive to changes in the flow rates, so the decrease in delta T would be greater at higher loads.

But the measured data is not behaving this way. Figure 9.95 shows the water side deltaT ($T_{wo} - T_{wi}$) vs load for the experiment week (red) and the week before it (blue). At low loads the delta T is actually greater during the experiment week than it was during the pre-experiment. And even at medium loads the delta T is equal in both cases, and the difference at high loads is not nearly as large as would be expected.

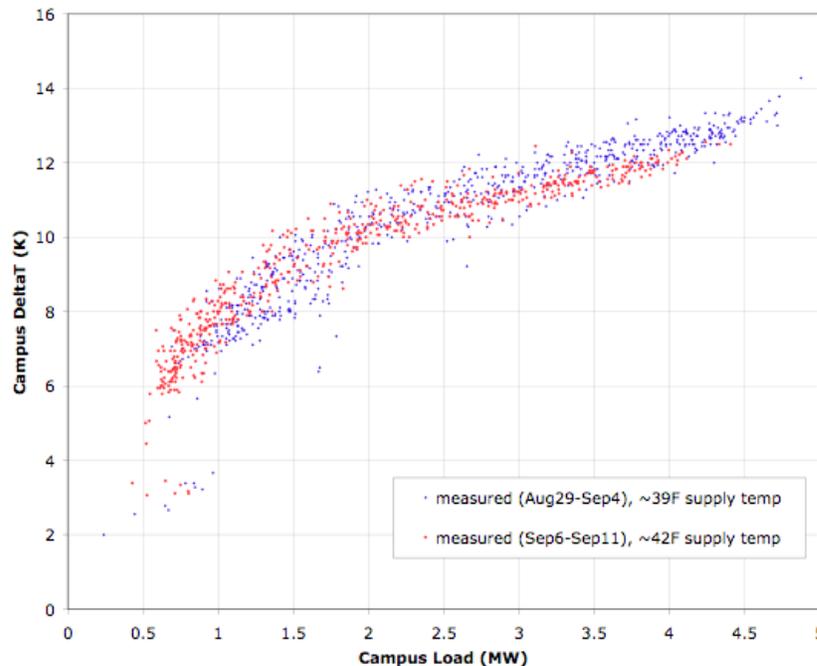


Figure 9.95. Water side delta T vs Load

By assuming values (as done in the model) for the return air temperature, supply air temperature and outside air fraction ($\gamma = 0.3$ in the figures below), along with the measured loads, outside air temperature and water side data, we can compute the LMTD and UA. (Note that although these are based on those three assumed values, the results are not very sensitive to them as long as they are within reasonable ranges.) The resulting graph of UA vs Load for the two weeks is shown in Figure 9.96. The offset is intriguing. And the graphs of UA vs water flow rate (Figure 9.97) and air flow vs water flow (Figure 9.98) are instructive - there seems to be an upwards offset in the graph of UA vs water flow during the experiment week, particularly at low loads, which does not seem to be explained by differences on the air side. This may be because of smoother control at the bridges and AHUs when the supply temperature is higher.

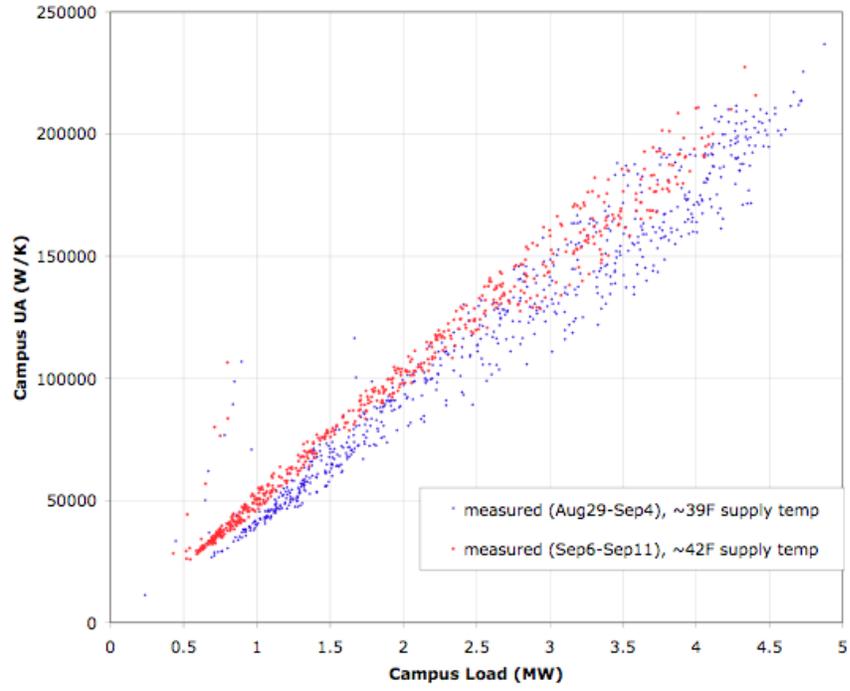


Figure 9.96. UA (calculated) vs Load

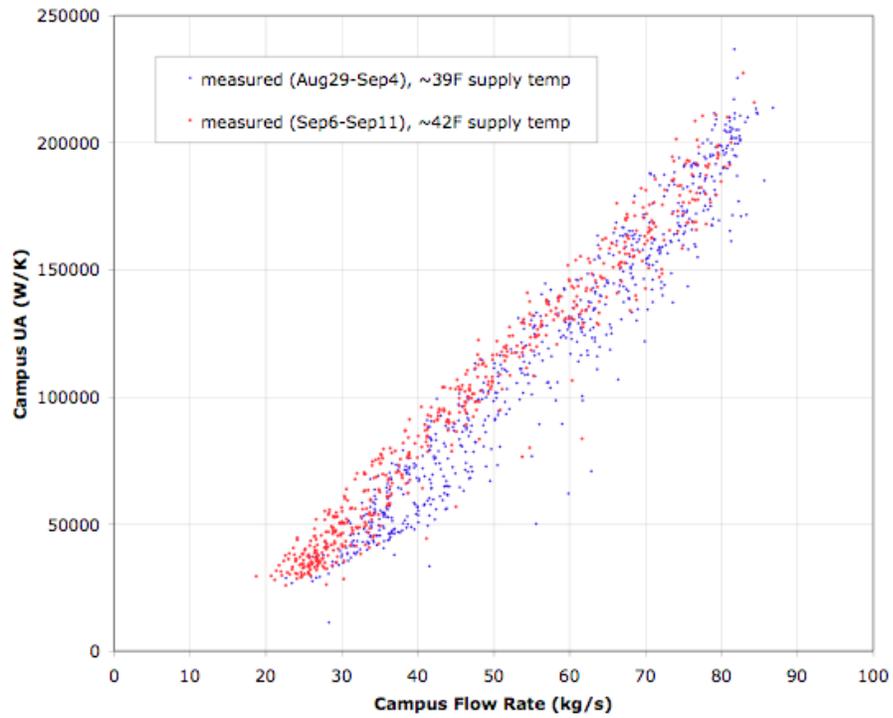


Figure 9.97. UA (calculated) vs Water Flow Rate

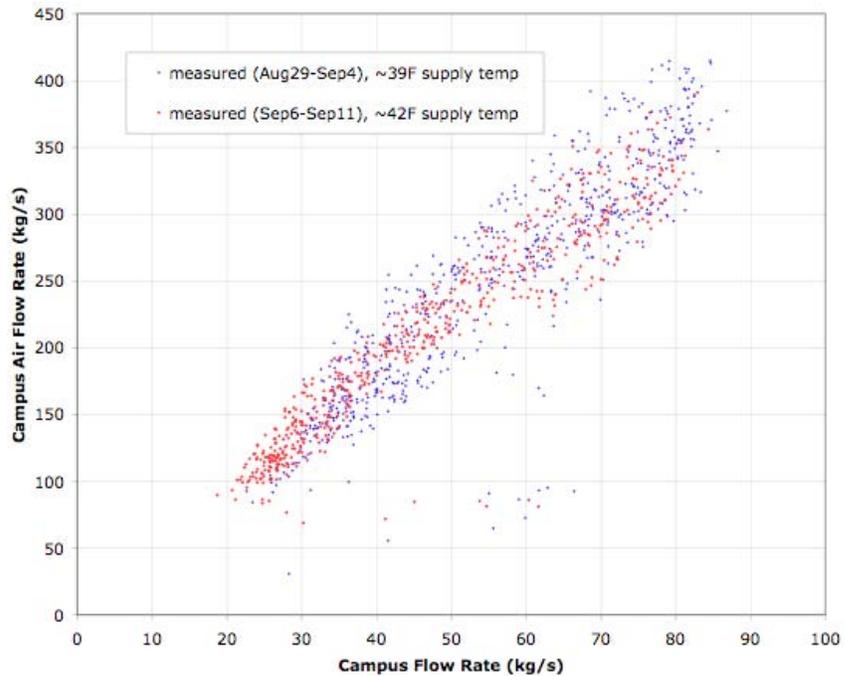


Figure 9.98. Air Flow Rate (calculated) vs Water Flow Rate

With this in mind, and for the purpose of having a functional empirical model in time for the second MPC experiment, an empirical factor was added to the model that scales the UA linearly with the supply water temperature between 3.9°C and 5.6°C (39-42°F), with the added model parameter c . With this added factor, the model is able to more accurately replicate the data,

as shown in Figure 9.99 and Figure 9.100.

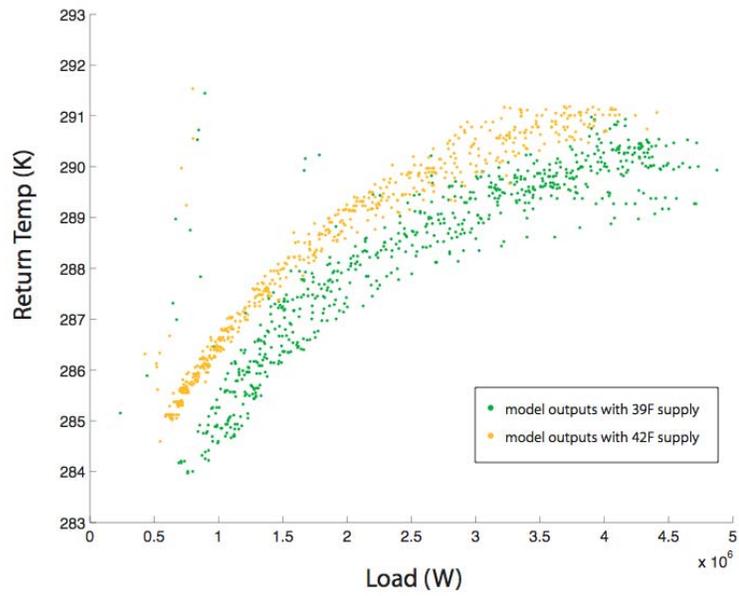
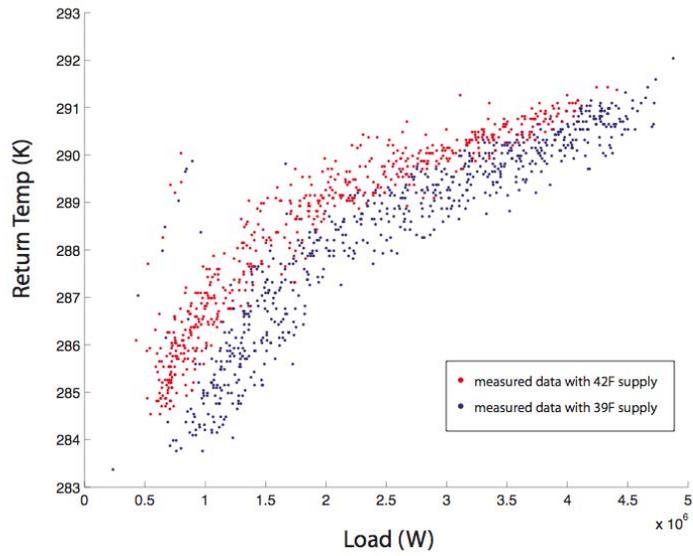


Figure 9.99. Measured data (top) vs model outputs (bottom)

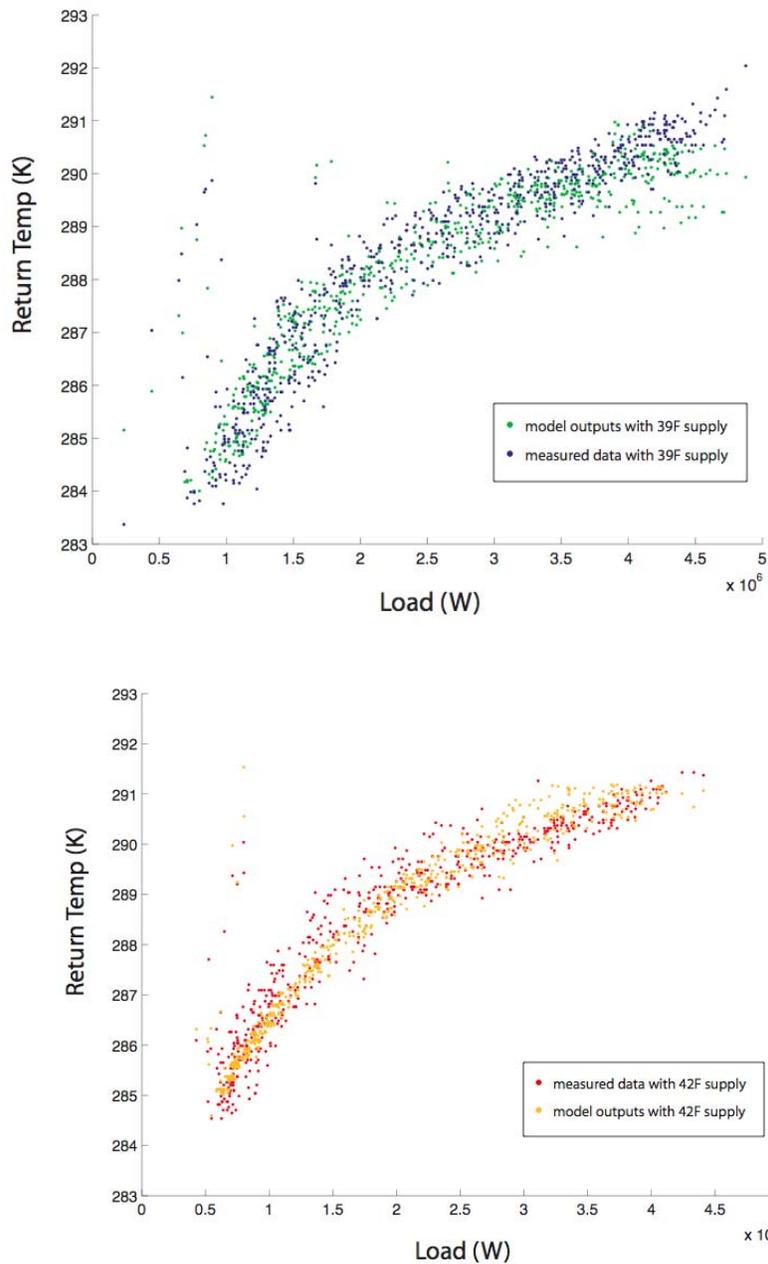


Figure 9.100. Measured vs model for standard supply temp (top) and higher temp (bottom)

This modified model was then used to re-calculate the lookup table for the chilled water return temperature model in the online model for the second MPC experiment. In the end, it did not affect the results of the MPC experiment because at almost all campus load levels, the decrease in campus delta T with the higher supply temperature had a greater negative impact on system COP than the positive impact resulting from the higher supply temperature, so the MPC was always pushing the supply water temperature towards its lower bound regardless.

9.4 Appendix 4: Development of an Implementation Based on Look-Up Tables

An implementation approach has been developed that is less intensive computationally than the approach used for the initial implementation and testing described in the rest of this report. The goal was to produce an implementation that could be used on a routine basis by the building operators without assistance from the research team and could also be used to reduce the computational effort required for the assessment of the performance and benefits of MPC in different scenarios, such as the building out of the campus. The implementation will be tested in follow-on work

9.4.1 Description of Approach

Figure 9.43 provides a highly abstracted view of the MPC implementation at UC Merced. This calculation is carried out every night at before 10pm, and the results are supplied to the operators for use that night.

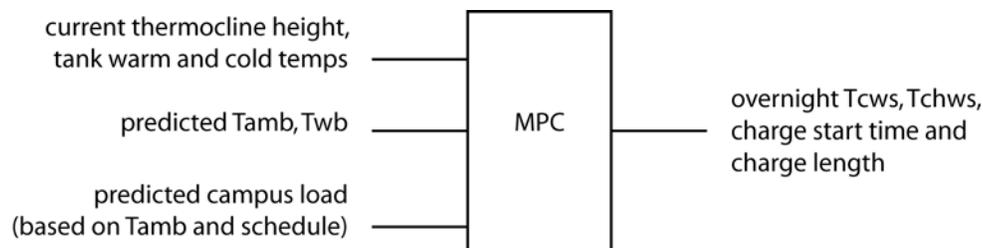


Figure 9.43: Top-level block diagram of MPC implementation

The experimental implementation performed this function each evening by solving a detailed optimization problem using the system model in Matlab and using Tomlab as the optimizer (a proprietary nonlinear optimization solver), which required about 20 minutes of computing time on a standard laptop. With this computing time requirement, an annual simulation of the MPC implementation would take approximately 5 days to run. In addition, the use of proprietary software makes it difficult to hand over a tool to the operators that they can use for long term implementation. The goal of the work described below is to decrease the online computation time requirement and to provide a program for long term implementation by simplifying the problem through some approximations and performing offline optimizations to produce a lookup table for a sub-problem, reducing the overall problem to a much simpler online optimization that can be solved quickly with non-proprietary solvers.

9.4.2 Problem Decomposition

The MPC system model (Figure 9.44) contains just three state variables: the tank thermocline height and the temperatures of the hot and cold nodes in the tank. The rest of the model is treated as pseudo-static. (The campus load model is dynamic, but it is external to the system model.)

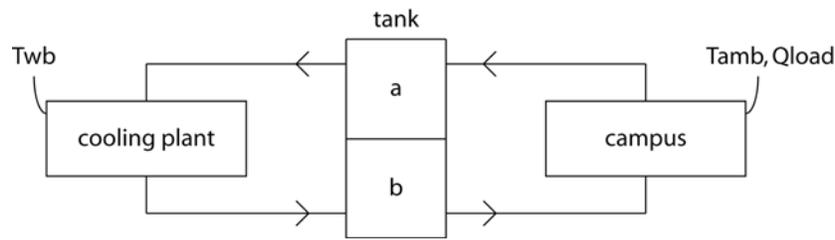


Figure 9.44: System model, overview

The MPC optimization problem can be rewritten in such a way to decouple the parts of the optimization that rely only on the cooling plant model from the parts that rely on the tank model as well. This allows for the offline creation of a lookup table of solutions to a sub-problem involving only the cooling plant model.

Consider the divisions of days and nights over the prediction horizon shown in Figure 9.45. Note that tank charging can only occur during the night period. This allows for a number of simplifications, such as the fact that the load from the campus during each day period can be treated as a single lump volume of water (m_{crd}) added to the tank between 8am and 10pm at a given average temperature (T_{crd}), and since the T_{wb} prediction is only required during the charging period, it only needs to be input for the overnight periods.

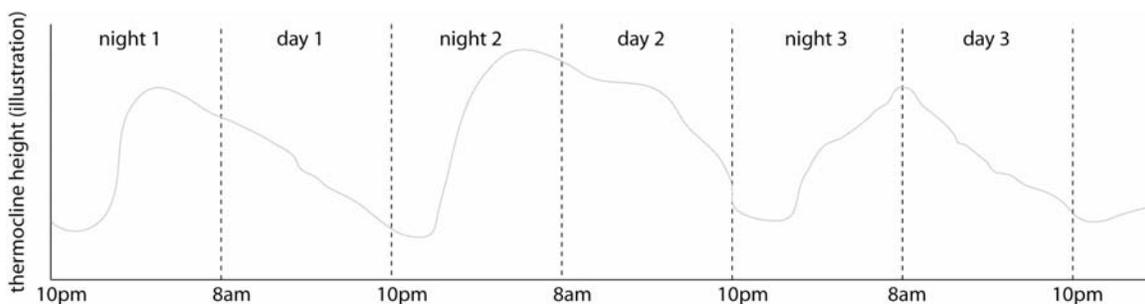


Figure 9.45: Days and nights over the prediction horizon

The MPC must determine the charge length for each night in the prediction horizon, the charging start time for each night, and the T_{cws} and T_{chws} for each night. Consider the following way of approaching this optimization:

- (1) A higher level optimization problem determines a charge length and T_{chws} for each night over the horizon
- (2) For each night, given the charge length and T_{chws} , a lower level optimization determines the optimal value of the start time and T_{cws} for that night, and returns the least-energy solution to the higher level algorithm.

Note that the solution of the sub-problem (Equation 9.4) for each night is independent of the solution for the other nights, and is dependent only on the given values of charge length and T_{chws} , the predicted overnight T_{wb} values, and the overnight values of the hot node in the tank. Since the solution of the higher level optimization problem depends on testing different values for the charge length and T_{chws} for each night, the sub-problem may be called many times in the course of solving the whole problem. The solution of this sub-problem requires only the cooling plant portion of the system model (Figure 9.44), which is fast running and free of state variables. This makes the sub-problem amenable to offline solution over a grid of possible conditions, resulting in a lookup table that may be used in place of online optimization.

The sub-problem is as shown in the following equation. Various approximations allow a lookup table of its solutions to be of a reasonably small number of dimensions.

$$[t_{start_i}^*, T_{chws_i}^*, Q_{plantConsumed_i}, T_{u8am_{i+1}}, T_{b8am_{i+1}}] = f(l_i, T_{chws_i}, \mathbf{T}_{wb_i}, T_{u10pm_i}, \mathbf{m}_{cRH_i}, \mathbf{T}_{cRH_i}, T_{b10pm_i}, h_{10pm_i}) \quad (9.4)$$

9.4.3 Input Approximations

The primary approximation that must be made is a parameterization of the predicted values of the ambient wet bulb temperature, T_{wb} . The experimental MPC implementation used 1-hr predictions. Using 1 hr predictions would require 10 dimensions in the lookup table for T_{wb} . This can be decreased by using just $T_{wb,min}$ and $T_{wb,range}$ as the basis for approximating the hourly predictions, or by using $T_{wb,min}$, $T_{wb,range}$ and the time of the min as the basis. For the first pass discussed below, the two-parameter approximation was used. The three-parameter version will likely be used for greater refinement.

The parameterization by $T_{wb,min}$ and $T_{wb,range}$ is described in Figures 9.46-9.48. Figure 9.46 shows the overnight T_{wb} values for a typical year. Figure 9.47 shows these same data with each night normalized by its minimum and maximum values: the colored lines are for each night, the dashed black line is the hourly average of the curves, and the solid black line is the normalized version of the dashed black line. This normalized curve is used to approximate the hourly T_{wb} predictions given the predicted min and range. Annual values of the min and range for a typical year are shown in Figure 9.49.

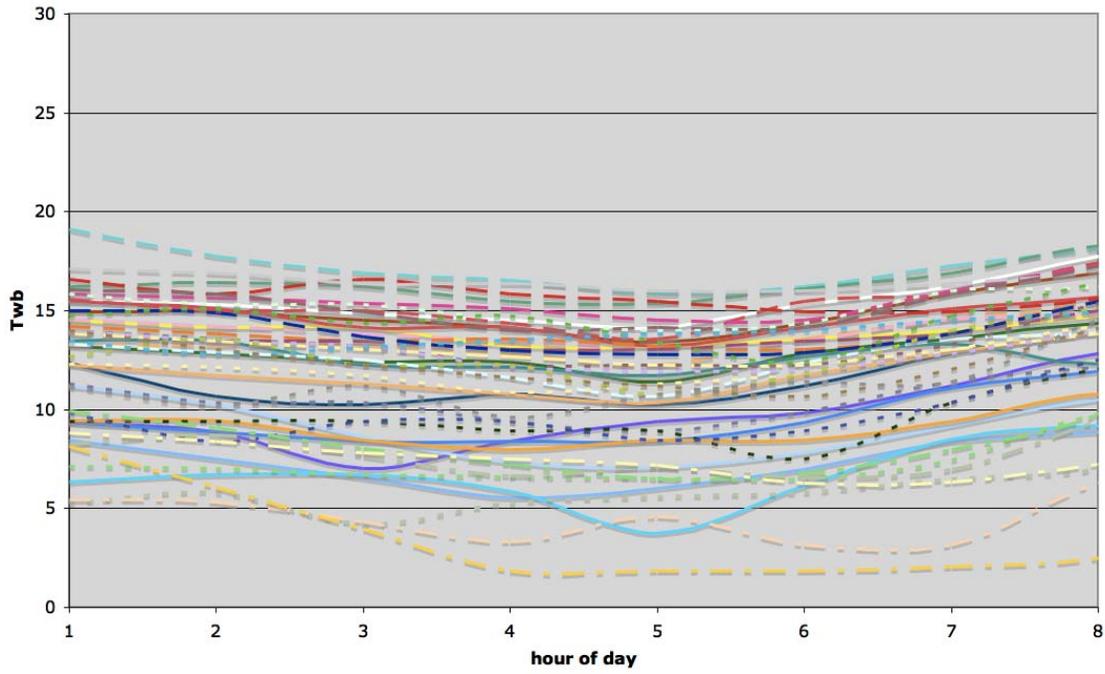


Figure 9.46: Overnight T_{wb} values, derived from EnergyPlus TMY file for Merced
 (Apr 1 to Nov 1, midnight to 8am, only 1 in 4 days shown for clarity)

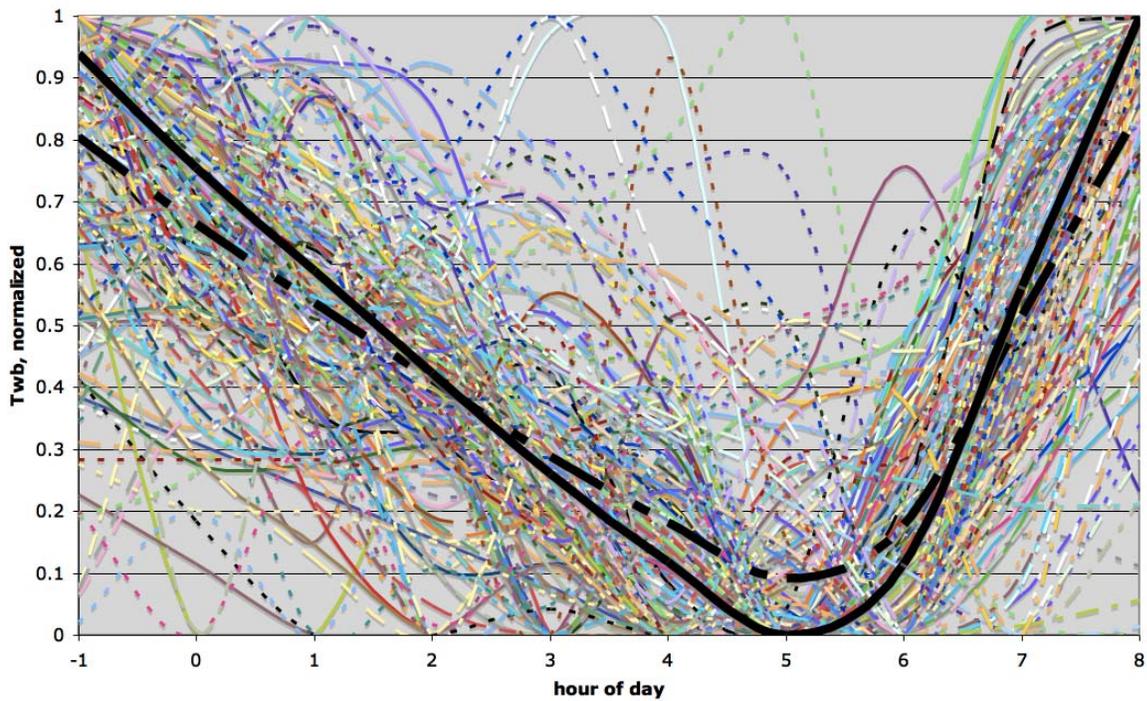


Figure 9.47: Normalized overnight T_{wb} values

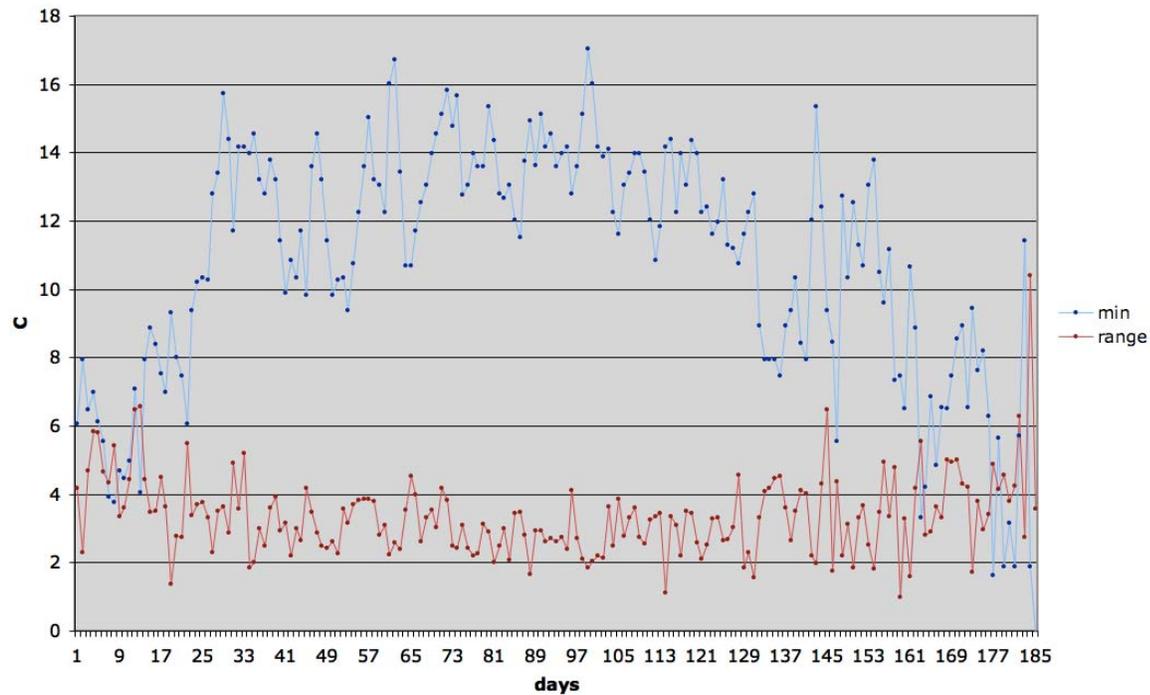


Figure 9.48: Overnight T_{wb} minimum and range over a typical summer (Apr 1 – Nov 1)

The other necessary approximations relate to the overnight temperatures of the hot and cold nodes in the tank. For the first pass discussed below, T_{chws} is assumed to be kept at a constant 39°F (as it was during the experiments). And, although it is a very rough approximation that will be the first to undergo refinement, the hot node in the tank (T_a) is assumed to be kept at a constant temperature. To capture the T_a temperature profile as accurately as the online MPC does, hourly predictions for the campus return flow and temperature would be required. However, it may be approximated by using an initial value for T_a and assuming a constant overnight flow rate (m_{crn}) and temperature (T_{crn}) from the campus. Or, a parameterization similar to that used for the T_{wb} prediction may be used. Measured campus flow rates for Sept 21 – Oct 7, 2009 are shown in Figure 9.49, and measured campus return temperatures for the same period are shown in Figure 9.50. Their influence on the temperature of the warm node in the tank is shown in Figure 9.51.

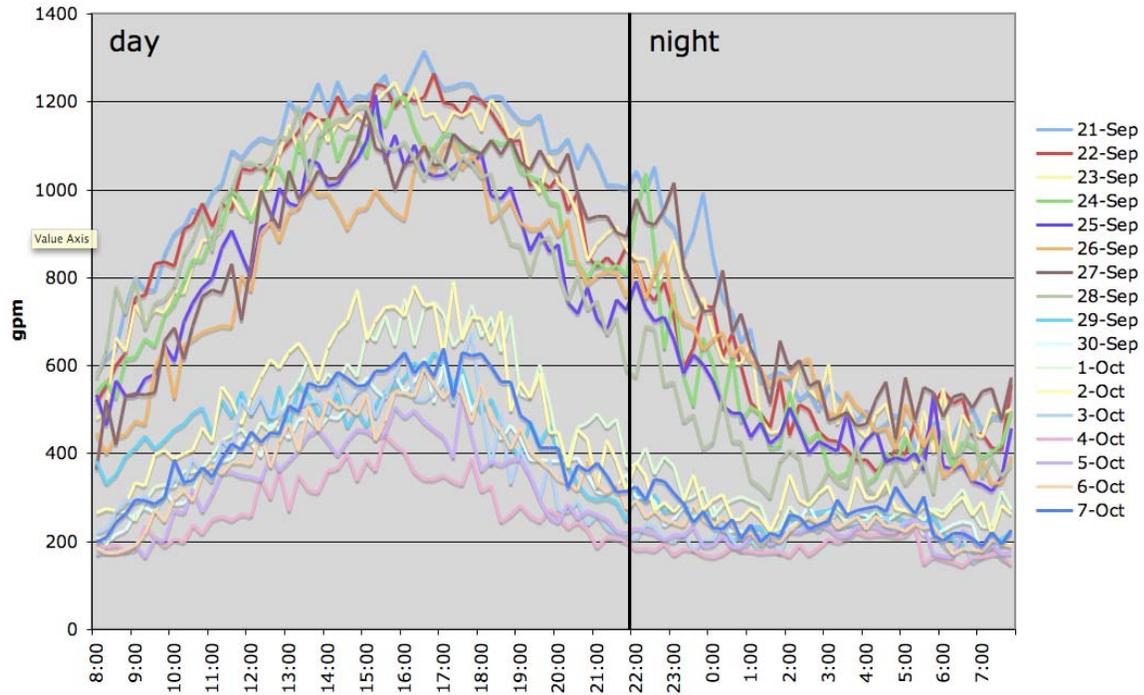


Figure 9.49: Measured campus return flow rates

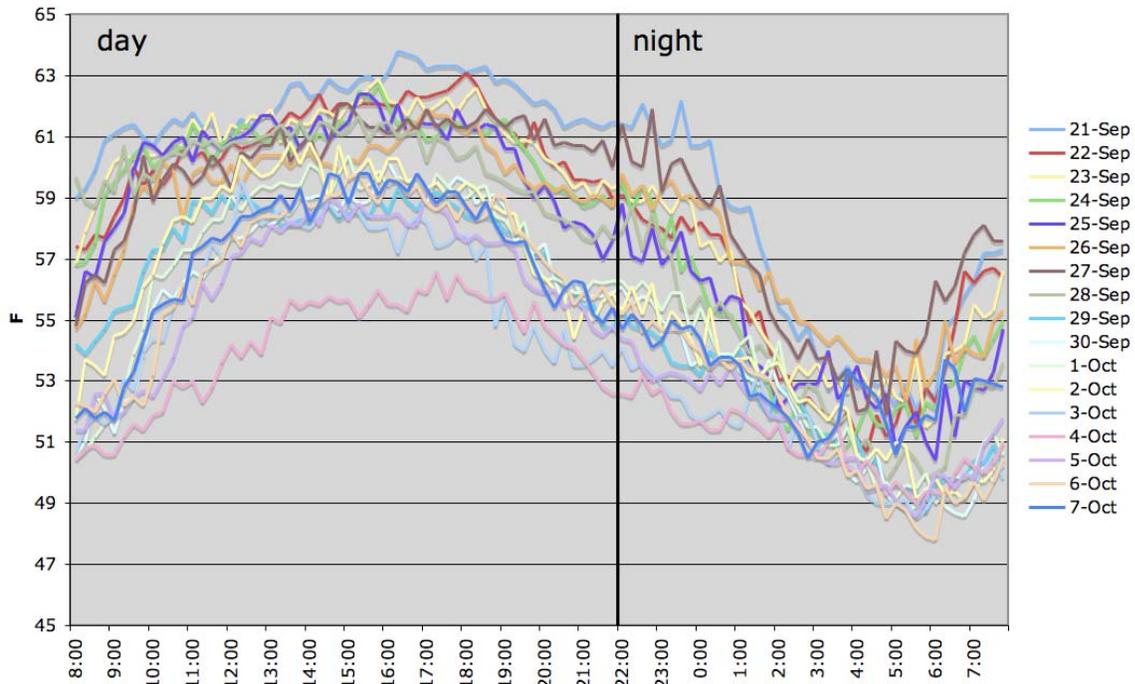


Figure 9.50: Measured campus return temperatures

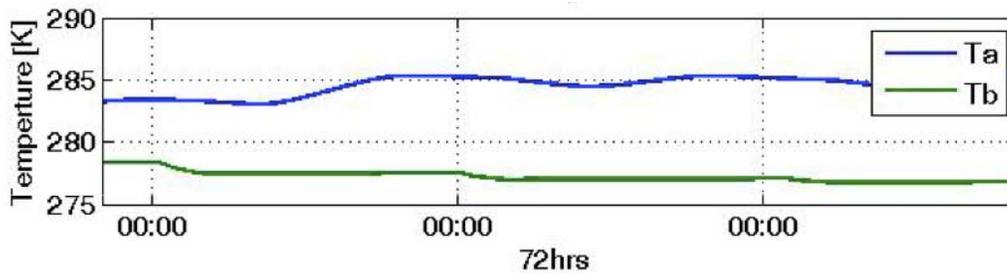


Figure 9.51: Tank temperature prediction (from one of the experiment #2 daily reports)

9.4.4 A First Pass at Approach Implementation – Lookup table creation

The approach outlined above is amenable to iteration, starting with rough approximations for the inputs and a relatively course gridding for the lookup table, then testing and refining the implementation over time. For a first pass, the two-parameter parameterization of the overnight T_{wb} prediction is used, along with the assumption of constant T_{chws} and T_{chwr} values. The original Modelica model of the chilled water plant was used for successive GenOpt optimizations over the grid of conditions shown in Table 9.4. The constraints on charging start time and T_{cws} in the Matlab implementation were also used in these optimizations.

Table 9.4: Grid of conditions for sub-problem solution lookup table

Parameter	Min	Max	Number of steps
Charge length (hrs)	1	10	10
$T_{wb,min}$ (F)	32	64.4	10
$T_{wb,range}$ (F)	1.8	18	10

Solving the sub-problem over this grid required 1000 optimizations, which took about 12 hours with one processor on a laptop. For each grid point, the solution vector contains a value for start time, T_{cws} and overnight energy consumption. Some slices through the resulting 3-dimensional conditions grids are shown in Figures 9.52-9.53. Figure 9.52 shows the optimal value of T_{cws} as a function of the $T_{wb,min}$. The scatter in this graph is because the optimal value of T_{cws} is also dependent on the values of $T_{wb,range}$ and the charge length. On the other hand, Figure 9.53 shows no scatter except a few outliers (perhaps because of local minima or the T_{cws} constraints?), suggesting that the charging start time is dependent only on the charge length. (It is also interesting to note that the relationship is very nearly linear.) This makes sense because our first pass T_{wb} prediction parameterization assumes the time of T_{cws} to be held constant, which will change in future iterations. Figure 9.54 shows another way of slicing the data for T_{cws} .

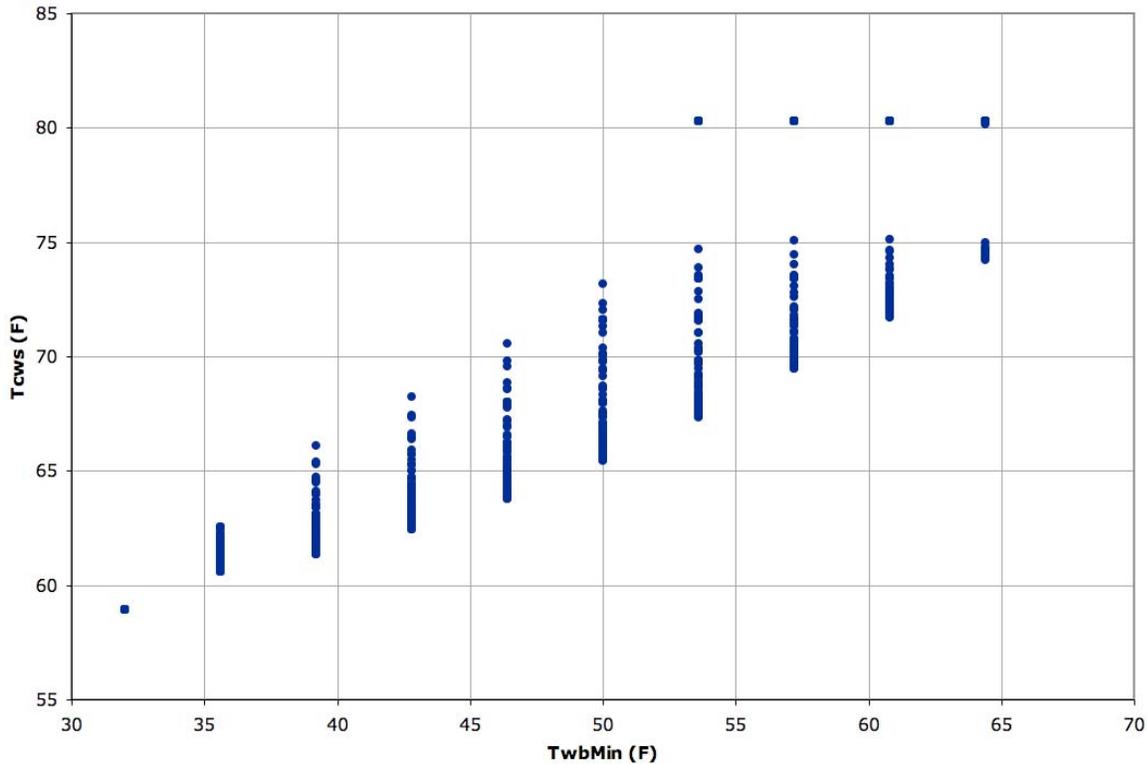


Figure 9.52: Optimal T_{cws} values vs minimum T_{wb} values

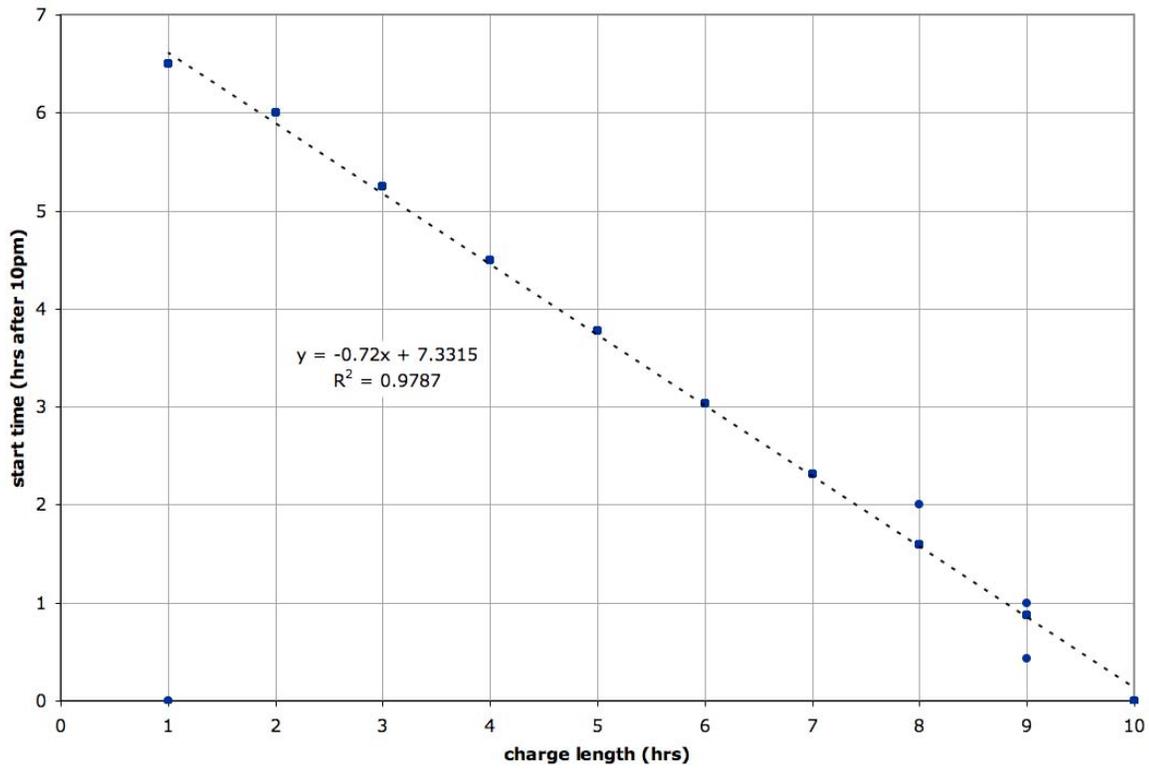


Figure 9.53: Optimal start time vs given charge length

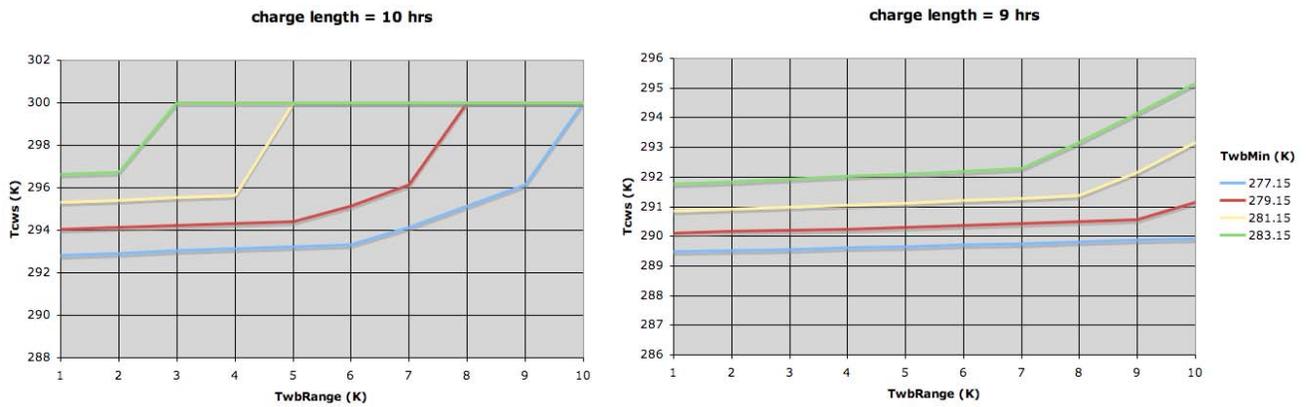


Figure 9.54: T_{cws} as a function of charge length, $T_{wb,Min}$ and $T_{wb,Range}$

9.4.5 First Pass at Approach Implementation – Online solution using lookup table

Given the lookup table of solutions to this sub-problem, the online MPC problem is simply to determine the charge length and T_{chws} for each night over the prediction horizon. (In this first pass, T_{chws} is held constant at 39°F, so the problem is just to optimize the charge lengths.) A horizon of 3 days (as used in the physical experiments) was used in this first pass, but this can easily be extended.

The objective function was coded in java, using linear interpolation over the lookup table. The table is currently small enough to be used in flat text file format, but SQLite is being considered for the next iteration. Penalty functions are used to enforce the thermocline height constraint. The java program takes as input a text file containing values of the following for each night over the horizon: the specified charge length, $T_{wb,min}$, $T_{wb,range}$, and the number of charge hours required to stay above the minimum thermocline height (which is calculated externally based on the initial thermocline height and the projected campus loads). The program outputs the total energy consumption plus any penalties for violating constraints, along with the optimal T_{cws} values and charging start times from the lookup table. This objective function program is called iteratively by GenOpt to find the optimal charge lengths.

This implementation has been tested using a couple of weeks of measured campus load and T_{amb} (T_{wb} approximated as $T_{amb} - 10$ K) from September 2009. The results are shown in Figure 9.55, with the charge length in red and T_{wb} in blue. Note that the approximate MPC algorithm is charging more on cooler nights, as expected.

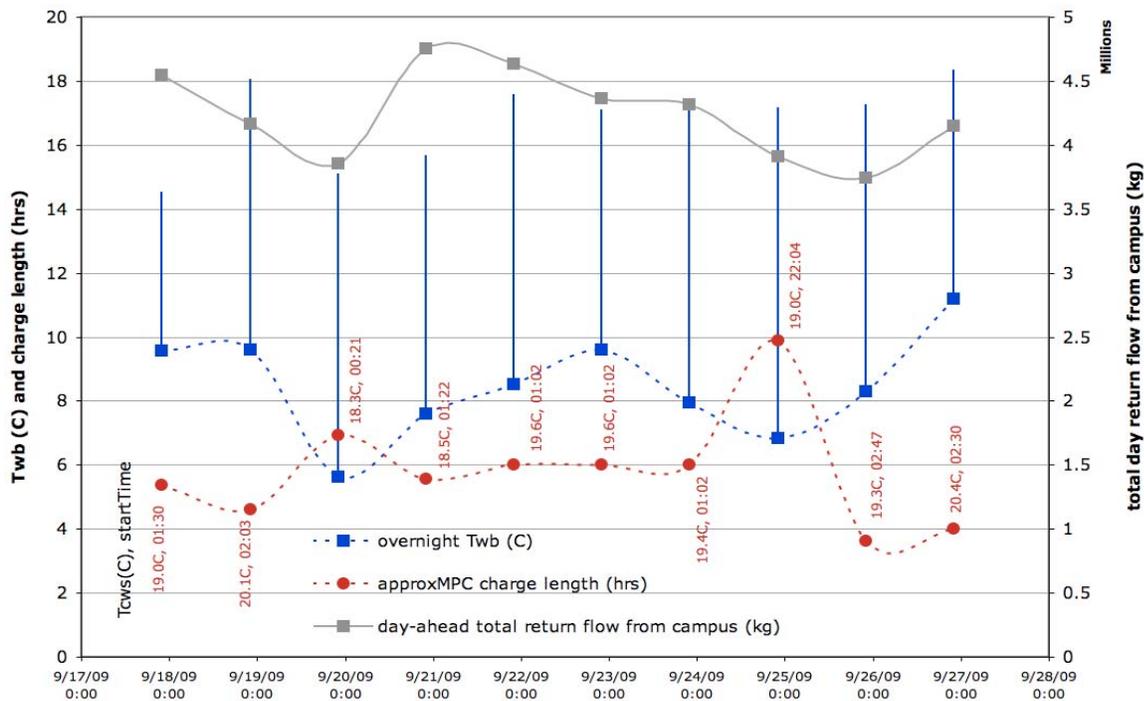


Figure 9.55: Results of first-pass test with online GenOpt solution using lookup table

Matlab MPC implementation), creating an interpolation lookup table equivalent to the Modelica model that can be run faster and without requiring Dymola licenses for multiple processors. The computation of the optimization lookup table could also be sped up by intelligently setting starting points for the optimizations based on previous solutions (e.g. start by solving for the corners of the lookup table, then interpolate to find optimization starting points for points in between).