NIST Technical Note 1707

Optimization of Building HVAC Systems Using Intelligent Agents – A Proof of Concept Study

Dr. George E. Kelly
Steven T. Bushby

NIST Technical Note 1707

Optimization of Building HVAC Systems Using Intelligent Agents – A Proof of Concept Study

Dr. George E. Kelly
Steven T. Bushby
.
Engineering Laboratory

October, 2011

U.S. Department of Commerce
Rebecca M. Blank, Acting Secretary

National Institute of Standards and Technology
Patrick D. Gallagher, Under Secretary for Standards and Technology and Director

National Institute of Standards and Technology Technical Note 1707
Natl. Inst. Stand. Technol. Tech. Note 1707, 21 Pages (October, 2011)
CODEN: NTNOEF

ABSTRACT

This paper describes the development and implementation of a simulation testbed used to study the concept of employing intelligent agents to optimize the performance of building Heating, Ventilating, and Air Conditioning (HVAC) systems. One day simulated cost savings results are presented for intelligent agents employing both a simple optimization method (SOM) and an advanced optimization method (AOM). The potential benefits and problems associated with each method are discussed and recommendations are made for future research efforts leading to the application of intelligent agents in real buildings.

INTRODUCTION

Over the last twenty years significant progress has been made in the integration of building control systems and building services through the development of standard communication protocols, such as BACnet and BACnet/IP (Bushby 1997). Unfortunately, little or no progress has been made in actually making "intelligent buildings" (also called "cybernetic buildings") intelligent or in optimizing the performance of building systems. Heating, Ventilating, and Air Conditioning (HVAC) control systems are still basically proportional–integral (PI) or proportional–integral–derivative (PID) at the lowest level, with one or two higher levels of heuristic supervisory control. Research on real-time optimal control was carried out by Z. Cumali (Cumali 1988), J. Braun (Braun 1989a, Braun 1989b), F. L. F. Pape (Pape 1990), J. House (House 1991), and others (ASHRAE 2007) in the late 1980's and early 1990's with mixed results. They showed that while optimization was possible, it was computationally intensive and difficult to implement in real building systems. This is because of the need to have information on the performance and status of all systems and equipment in one place and the difficulties in handling boundary conditions and the discontinuous systems found in buildings (e.g., HVAC equipment that turn on and off, systems that change operating states, systems with multiple modes of operation).

The use of artificial intelligence (AI) techniques (Barr 1981) and, in particular, intelligent agents offers a possible solution to these problems. The latter has been successfully implemented in a variety of applications (Barr 1982), including search engines and robotic systems, and a considerable amount of information already exists in the AI community on different agent architectures (e.g., deliberating, reactive, and hybrid), agent design and implementation, and agent programming. Intelligent agents know or can learn the performance and status of the systems and equipment they monitor and can communicate and collaborate with other agents to achieve a common goal, such as minimizing energy consumption and/or cost of operation, maximizing comfort, identifying and diagnosing problems. Intelligent agents make it possible to solve the problem of building system optimization in a distributed manner which greatly simplifies the computational methods required.

To study the application of intelligent agents to the optimization of HVAC system performance, a simulation testbed was developed and used to study how intelligent agents can learn the performance (referred-to in this paper as "identification" (Ljung 1987)) of the building systems and equipment they monitor and then communicate and collaborate with other agents to achieve a common goal, such as minimizing energy consumption or cost of operation. Two optimization methods were developed and evaluated, a simple method and an advanced method. The following sections of this paper discuss the testbed, the identification or learning method employed, the simple and advanced optimization methods, some cost savings results, and conclusions and recommendations.

It is important to note that while the name “intelligent agents” is used in this paper for the program modules in the testbed that perform identification and optimization, they are far from being intelligent. In fact, they are rather “simple” in comparison with “intelligent agents” used in other applications. However, this is intended to be a proof of concept study and, as with any new research area, it is often best to start out simple and then move on to more complex approaches. In addition, it has been shown again and again that “simple” is often “best” when trying to do control or perform fault detection or commissioning on HVAC systems, and the optimization of such systems is not likely to be any different.

THE TESTBED

The testbed currently consists of 17 “agents” written in the Java programming language running in a free framework program (Pantic 2005) that was developed at the Delft University of Technology as an AI teaching tool. The agents simulate the performance of different building HVAC system components and exchange information using publish/subscribe “channels” within the framework. There are eight variable air box (VAVBox) agents that simulate VAVBox/zone interactions, two air handling units (AHUs) which include cooling coils and variable speed supply air fans, two variable speed chillers, two variable air flow cooling towers, an agent for coordinating the control actions of the HVAC agents, an agent for accumulating and reporting costs, and an agent for controlling the simulation (SimulationTimer Agent).

A block diagram of the fourteen agents that simulate various building/HVAC components (i.e., VAVBoxs and cooling loads, AHUs, chillers, and cooling towers) is shown in Figure 1a. They all have a similar structure. They first import the various Java packages needed to run the agent and then set up publish and subscribe channels that the agents use to communicate with each other. Initial values are then defined and message “handlers” are setup to receive and parse messages on each subscribe channel. Whenever a message is received it is immediately processed.

The SimulationTimer agent, which gets the “time” from the computer running the testbed, sends the time to the other agents every 5 seconds. Each time an HVAC agent receives the time it runs that component’s simulation. In the current implementation of the testbed, a second in the simulation corresponds to one minute in “real” time. Thus a 12 hour “day” would take 12 minutes to simulate. After running each component simulation (every 5 seconds), the updated variable information (e.g., the air flow rate to each zone) is sent to other agents who use the information to carry out their own component simulation. A single “SimulationData” channel is used to both publish and receive information need to carry out the simulation.

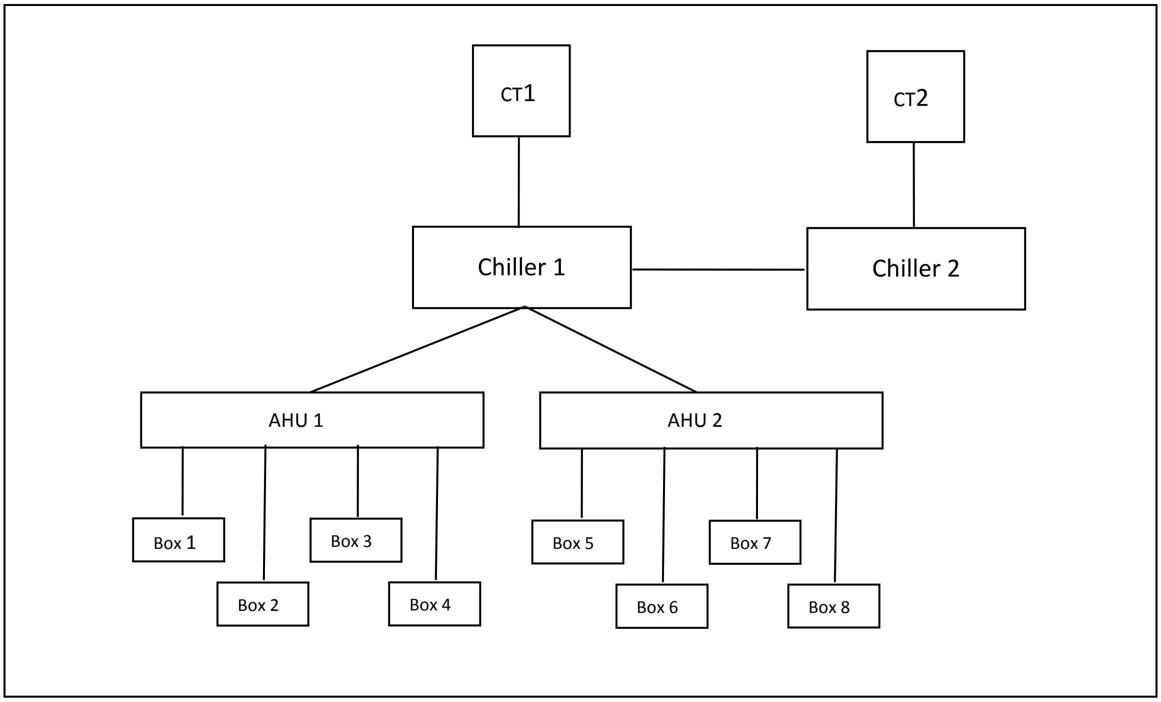

Figure 1a. Diagram of HVAC simulation agents in the testbed, including Cooling Towers CT 1 and CT2, Chillers 1 and 2, AHUs 1 and 2, and VAV Boxes 1 through 8.

In the current testbed, small time delays are employed in the simulation so that each component type (i.e., VAV boxes, AHUs, chillers, and cooling towers) carries out its particular simulation at a particular point in the 5 second simulation period. Thus, all the VAVBox/zone simulations are first performed and the updated information sent out on the SimulationData channel. This information is then read by the AHU agents, which then perform their own component simulations. They then send out their updated information on the SimulationData channel for use by the chiller agents. The cooling tower agents then take their turn receiving information, simulating the performance of the cooling towers, and sending updated information out on the SimulationData channel. Finally the "Coordinator" agent and a "CostMeter" agent run. The former coordinates the control actions of the HVAC agents. The latter, which has been collecting "cost rate data" (i.e., cost per hour information) sent to it by the other agents during this 5 second simulation period, sends the information out on a "PlotData" channel in a format that can be easily processed and plotted using a commercially available spreadsheet. While these delays may or may not be necessary from a simulation point of view, without them it would be extremely difficult to debug the various HVAC simulation agents in the testbed because information sent by an agent at one simulation time step could end up being used by another agent in a different time step.

The first time an AHU, a chiller, or a cooling tower simulation agent receives the time from the SimulationTimer agent it setups a "timer" to periodically run an "intelligent agent", which will henceforth be referred to in this paper as an "Agent" (with a capital A), that does the identification and optimization. This Agent runs independently of the simulation agent. It receives information from the simulation agent through shared variables, and exchanges information with other Agents over the same SimulationData channel used to exchange simulation information. Although the testbed was designed to allow the various Agents to run at timesteps that were different from the timestep used by the simulation agents (currently 5 seconds or 5 minutes in real time), this was not implemented in the current version. Instead, the "timers" that trigger the running of the different Agents were set so that each Agent ran in the same time interval (described above) within the 5 second timestep used by the simulation agent that created it. This was done to simplify debugging and because it seemed best to have the Agents run as frequently as possible (i.e., with the smallest possible timestep) when doing optimization, but not more frequently than the simulation agents.

Figure 1b shows how information is exchanged between two simulation agents and the Agents they create.

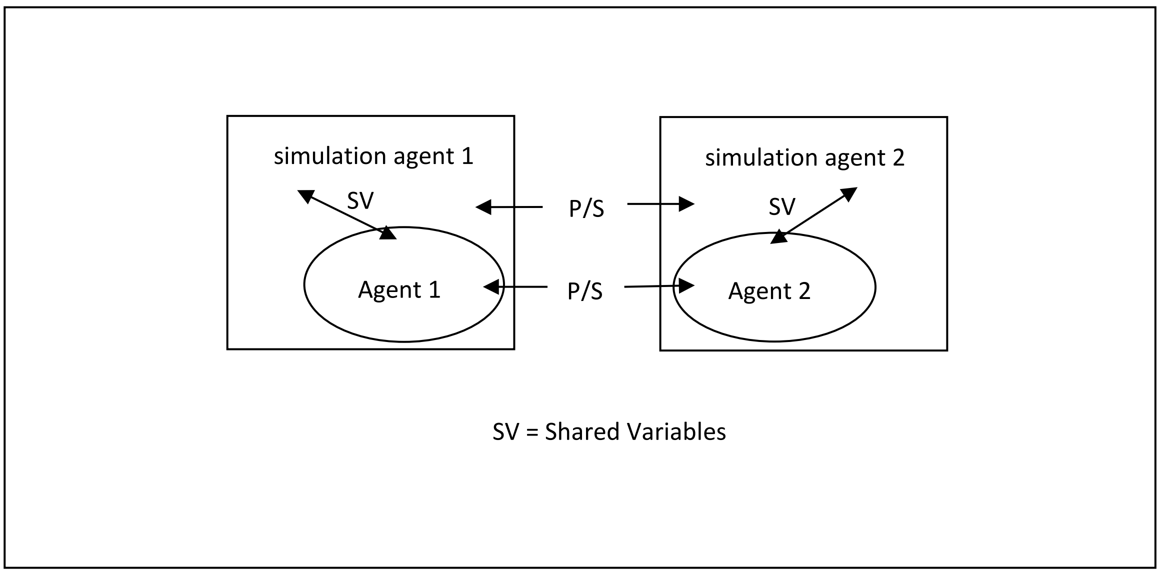

Figure 1b. Information exchange between simulation agents and the Agents they create. The Agents perform identification and optimization.

HVAC SIMULATION MODELS

The VAV box model uses a simple energy balance to determine the zone air temperature and air flow rate into and out of the zone. A 10 minute time constant is used to simulate the zone dynamics. External loads follow a half-sine function with a peak load at 2:00 PM and zero loads before 8:00 AM and after 8:00 PM. Internal loads are varied each hour between 6:00 AM and 5:00 PM in a manner that simulates different "typical" zone occupancy patterns.

AHU1 handles VAV boxes 11, 12, 13, and 14; AHU2 handles VAV boxes 21, 22, 23, and 24. The cooling coil model is a second order polynomial that relates the water-side effectiveness to the fraction of mass water flow rate for a given air flow. It calculates the total supply/return air flow rate and the mixed return air temperatures. This, along with the supply air temperature and the chilled water supply temperature, are used by the cooling coil model in the AHU agents to calculate the chilled water flow rate and chilled water return temperature.

The chiller model, which was obtained from a leading chiller manufacturer, is a “generic model” of a variable speed centrifugal chiller that was developed for ASHRAE Standard 90.1. It uses three second order polynomials to relate capacity, power, and part load performance to three independent variables, chilled water supply temperature, chiller part load ratio, and entering condenser water temperature. Power Ratio curves (i.e., curves of power/maximum power at standard rating conditions) for this chiller model are shown in Figure 2 for an entering condenser water temperature of 23.3 ˚C (74.0 ˚F) and part load ratios (plr), defined as chiller capacity divided by the maximum chiller capacity at reference conditions, of 0.9, 0.8, 0.6, and 0.4.

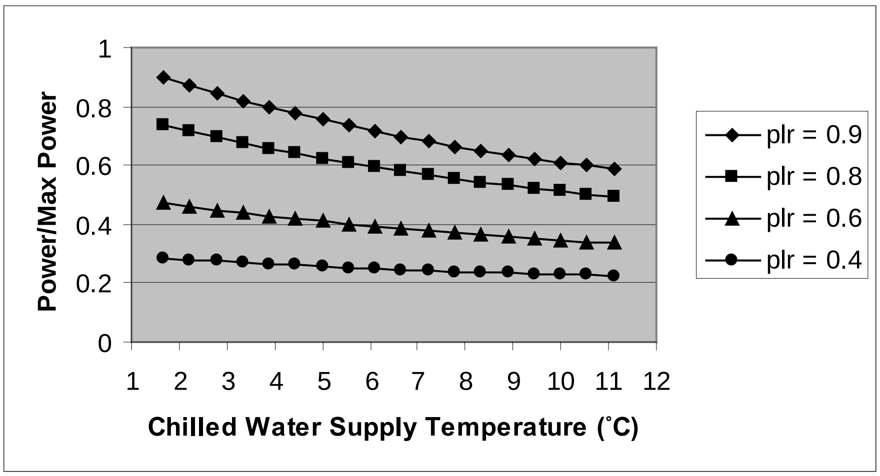

Figure 2. Chiller Power Ratio Curves for part load ratios of 0.9, 0.8, 0.6, and 0.4 and an entering condenser water temperature of 23.3 ˚C (74.0 ˚F).

Chiller1 is considered to be the “lead” chiller. Its agent calculates the temperature of the chiller water returning from the AHUs, decides how much of the total cooling load will be handled by each chiller, and sends the information to Chiller2. Each chiller calculates the temperature of the condenser water leaving the chiller. There is a one-time step delay of 5 seconds (5 minutes in “real” time) between the Chillers receiving information from the AHU agents and using it to simulate chiller performance. Although this is an artifact of the simulation methodology employed, it can be considered to correspond to the “dead time” required for the chilled water to return from the AHUs to the chillers.

The cooling tower model employs a fifth order polynomial with twenty-seven coefficients that relates the Approach (defined as the temperature difference between the water temperature leaving the cooling tower and the ambient wet bulb temperature) to the air flow rate ratio (defined as the ratio of the air flow rate to the maximum air flow rate), the water flow rate ratio (defined as the ratio of the water flow rate to the maximum water flow rate) , the range temperature (defined as the inlet water temperature minus the outlet water temperature), and the inlet air wet-bulb temperature. CoolingTower1 serves Chiller1, while CoolingTower2 serves

Chiller2. Each cooling tower model calculates the air flow through the tower that will give a specified Approach. In the current version of the testbed, the condenser water flow rate is held constant.

The chilled water and condenser water pump models and the AHU fan models use second and third order polynomials to calculate the "variable speed" performance costs of pumps and fans. A "work day" during which energy consumption data is collected runs from 7:00 AM to 5:00 PM.

IDENTIFICATION

In a real building/HVAC system, an intelligent agent must first learn (identify) the characteristics of the equipment it is responsible for before it can attempt to optimize its performance. The existence of nonlinearities in these systems and, in particular, the models described above make on-line identification difficult. To develop the best identification method for implementation in the testbed, the HVAC component models described above were exercised over the allowable range of independent variables and the results were compared with different simpler models. Both multivariable linear regression and piece-wise linear regression in one or two variables were examined to determine which simplified model worked best in each application.

For the cooling coil, piecemeal linear regression was used to relate the water side effectiveness to (1) the ratio of the water flow rate through the coil to the maximum water flow rate and (2) the ratio of air flow rate through the coil to the maximum air flow rate. However, while this works reasonably well, it was not used in the intelligent agent optimization process because the chiller model did not use the returning chilled water temperature and the energy consumed by the chilled water pump was small in comparison to that used by the chiller. This allowed the effect of changes in chiller water return temperature and chilled water flow rate to be ignored in determining the optimal supply air and chilled water temperatures setpoints. It also allowed for a simpler approach to be employed in carrying out joint chiller-AHU optimization that involved not permitting the difference between the supply air temperature setpoint and the chilled water temperature setpoint to be less than 5.56 ˚C (10.0 ˚F). This simplification is based upon a separate analysis of a typical cooling coil, which showed that (1) for a properly sized coil and large cooling loads, a supply air temperature set point 5.56 ˚C (10.0 ˚F) above the chilled water temperature setpoint was either optimal or very close to optimal, and (2) for moderate and low cooling loads, while the optimal supply air temperature setpoint was closer to the chilled water temperature setpoint, the effect on the cost of operation of keeping it 5.56 ˚C (10.0 ˚F) above the chilled water temperature was very small. (See Appendix A.)

For chiller performance identification, all three independent variables (i.e., the chilled water supply temperature, the chiller part load ratio, and the entering condenser water temperature) were required to do the performance identification. The part load performance curves giving chiller power as a function of the variable x , defined as chiller capacity/maximum chiller capacity at reference conditions, were determined over two different load ranges ($0.3 < x <= 0.6$ and $0.6 < x <= 0.9$) for two different values of chilled water supply temperatures, 2.78 ˚C (37.0 ˚F) and 6.67 ˚C (44.0 ˚F), and two different values of the entering condenser water temperature, 23.3 ˚C (74.0 ˚F) and 26.9 ˚C (84.0 ˚F). This created eight regions for which the part load performance curves were determined (by linear regression) as a linear function of x. These curves were then used to calculate the chiller's power ratio, pwrRatio, for eight chilled water supply temperatures between 0.56 ˚C (33.0 ˚F) and 12.2 ˚C (54.0 ˚F)) and the current value of entering condenser water temperature. The chiller power consumption at a specific chilled water supply temperature was found by interpolation between these eight values.

The above chiller performance identification process required two simulation days (Day #1 and Day #2) to carry out and gave reasonably good results as shown in Figure 3.

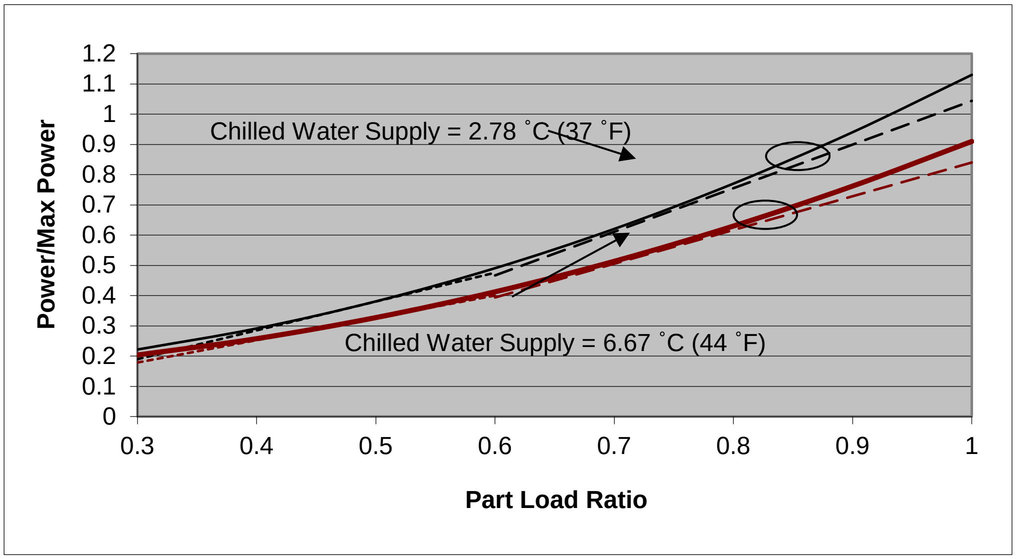

Figure 3. Comparison of the power consumption (shown as a power ratio) of the identified chiller model with the actual chiller model for an entering condenser water temperature of 25.6 ˚C (78.0 ˚F) and two different chilled water supply temperatures. The "solid lines are results from the Chiller simulation model; the dotted and dashed lines show results from the piece-wise linear identification process.

In developing a performance identification method for the cooling towers, it was found by trial and error that the number of independent variables could be reduced by defining the new variable, w, given by:

$$w = \frac{\text{Approach}}{\text{Trange} + 1.67\ ^\circ\text{C}}$$

where Trange is the cooling tower range temperature (˚C).

For different wet-bulb temperatures, one finds that the results obtained for the cooling tower Liquid-Gas ratio (LG ratio), defined as the ratio of the water flow rate ratio to the air flow rate ratio, are clustered together and can be approximated by straight lines over a wide range of the variable w. This is illustrated in the Figure 4 for a wet-bulb temperatures of 20 ˚C (68 ˚F) and

Approaches of 5.56 ˚C (10.0 ˚F), 11.1 ˚C (20.0 ˚F), and 16.7 ˚C (30.0 ˚F). The thicker line in this figure illustrates that a linear approximation to the data works well. Similar results were obtained at other wet bulb temperatures.

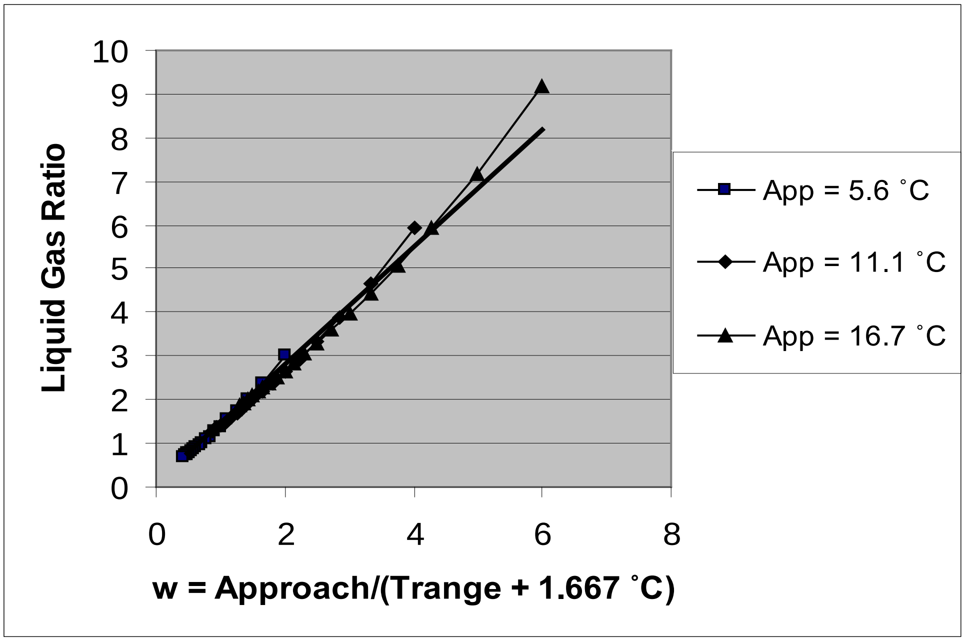

Figure 4. Plot of the Liquid-Gas ratio predicted by the cooling tower model used in the HVAC system simulation against the variable w = Approach/(Trange + 1.67 ˚C). Results are for a wet bulb temperature of 20 ˚C (68 ˚F); Trange was varied between 1.1 ˚C (2.0 ˚F) and 11.1 ˚C (20.0 ˚F). The thick line illustrates a linear approximation to all the data shown.

To determine whether this is an adequate simplification for identifying cooling tower performance, the testbed was run over a period of one day (from 6:00 AM to 6:00 PM) for a wet-bulb temperature of 20 ˚C (68 ˚F) and an Approach of 5.56 ˚C (10.0 ˚ F). Linear regression was then used to find the LG Ratio as a function of the variable w. The results are shown in Figure 5. Series 1 is the actual results from the simulation model. Series 2 is the predicted results using the equation obtained by linear regression. The top graph gives the “actual” and predicted LG Ratio for each time step in the one day simulation, while the bottom graph plots the “actual” and predicted LG Ratio against the variable “w”. Similar results were found for approaches equal to 11.1 ˚C (20.0 ˚F) and 16.7 ˚C (30.0 ˚F). The excellent agreement between the “actual” and predicted results indicates that this simplified method of identifying cooling tower performance appears to work well.

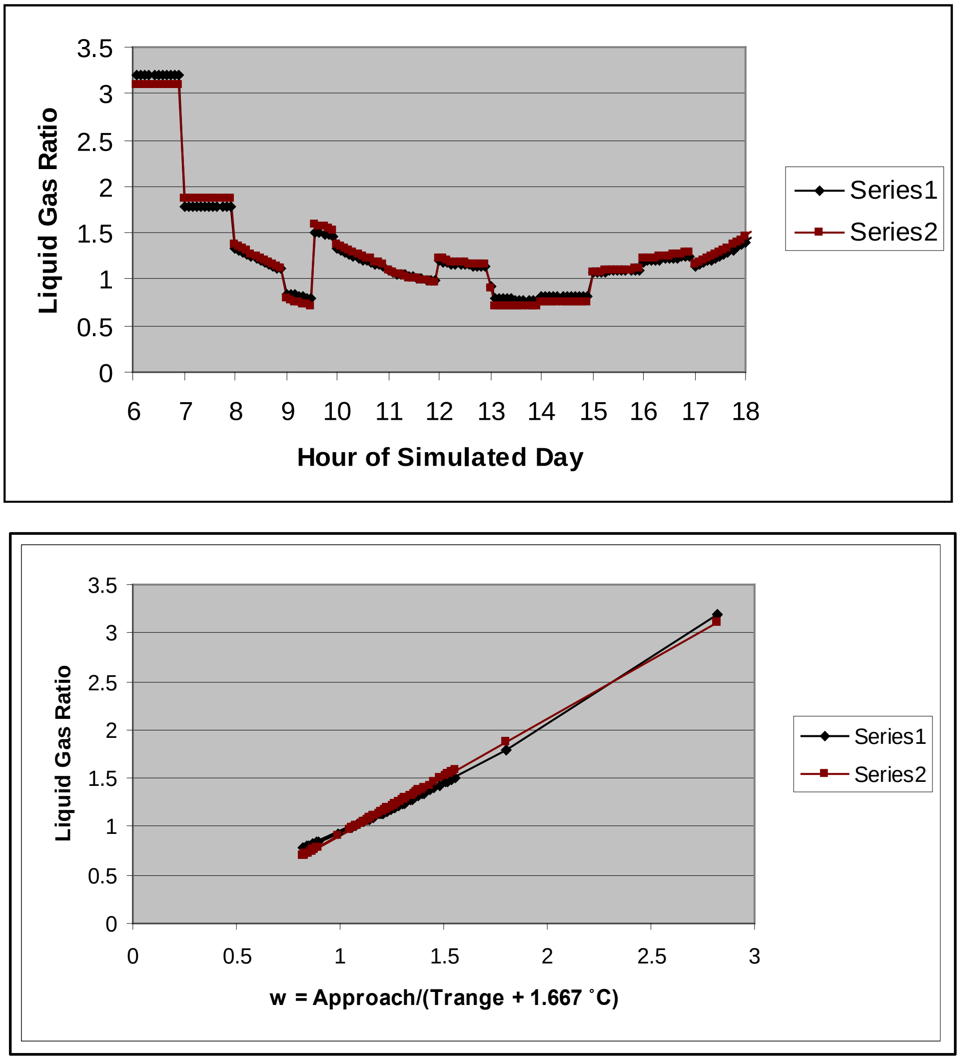

Figure 5. Comparison of the Liquid-Gas ratio predicted by the actual cooling tower model (Series 1) with that predicted the simplified model (Series 2). The top graph shows the results for both models over a simulated day. The bottom graph shows the same results plotted against the variable w = Approach/(Trange + 1.67 ˚C).

SIMPLE OPTIMIZATION METHOD

A simple optimization method (SOM) was the first method implemented in the testbed. In the SOM, the various Agents (after identifying the performance of each piece of HVAC equipment) take turns optimizing the performance of the particular piece of equipment for which they are responsible. To this end, each of these Agents was given the ability to evaluate the effect of a proposed set point change on the "rate of operating cost" (i.e., cost/hour) of the particular HVAC component (e.g., chiller) it was responsible for and to communicate with other Agents to determine what effect such a change would have on the "rate of operating cost" of other HVAC components (e.g., AHUs). After the total cost of making the proposed set point change has been determined, the Agent proposing the change decides whether or not to proceed with the set point change. Each Agent participating in the optimization process takes turns adjusting the set points under its control to minimize the cost of operation of the building system.

One approach to implementing the SOM is to employ a "token" that is passed from one Agent to another. When an Agent receives the token it decides, based upon cost information it calculates for the piece of HVAC equipment it controls and cost information it receives from other Agents, whether to raise or lower or keep unchanged its particular setpoint. It then passes the token on to the next Agent which then adjusts its own setpoint. While this type of token passing might be ideal for a real building HVAC control system, it causes problems in a research project where it is important to examine the effect of various combinations of Agents working together to implement a simplified optimization method. Thus, in the current version of the testbed it was decided to employ a Coordinator Agent. This Agent's job is to dictate which Agents are allowed to change their setpoints and the order in which the Agents act. This has the same effect as token passing but avoids the difficulty of having to rewrite computer code every time the number of Agents employed in the optimization process changes. It also simplified the development of an Advanced Optimization Method (discussed below) since such a Coordinator Agent would be able to receive "optimization proposals" from different Agents and decide which one should be implemented. The alternative of having the different Agents negotiate among themselves, while probably better in the long run, was beyond the scope and intent of the current research effort.

As implemented in the SOM, the Coordinator Agent sends a "take action" token to a particular HVAC Agent that tells the Agent that it may now raise or lower or keep the same the setpoint it is responsible for. When the HVAC Agent has completed its "action", it sends a "task completed" token back. When the Coordinator receives this "task completed" token, it can either send a new "take action" token back to the first HVAC Agent or to another HVAC Agent. By simply changing the token the Coordinator Agent sends to each HVAC Agent, the Coordinator can control which HVAC Agents are involved in the optimization process. Changing the tokens in the Coordinator involves only changing a few numbers and is simple to implement.

After an HVAC Agent receives a "take action" token from the Coordinator Agent, the HVAC Agent uses the identified simple model to calculate the "rate of operating cost" (in dollars per hour) of the HVAC equipment it is responsible for at the current setpoint and at two "proposed" new setpoints, one above and one below the currrent setpoint. For the AHU and Chiller1 Agents, 1.1 ˚C (2.0 ˚F) above and below the current supply air temperature setpoint and the current chilled water temperature setpoint, respectively, are used as proposed new setpoints. For the cooling towers, 1.1 ˚C (2.0 ˚F) above and below the current Approach setpoint are used. The HVAC Agent also sends an "inquireCost" message to the other affected equipment's Agent asking for information on the rate of operating cost at conditions corresponding to the current setpoint and the two proposed setpoints. A receiving Agent then uses information obtained from its performance identification process to calculate and then return the requested information. The Agent who received the "take action" token then calculates the "total" change in the rate of operating cost at the two proposed setpoints. If the total rate of operating cost can be reduced by changing the current setpoint to one of the proposed setpoints, it sends a "requestAction" message to itself. If this (proposed) setpoint change also involves another Agent changing its setpoint, the Agent receiving the "take action" token also sends out a "requestSetpointChange" message. The Agent receiving the "requestSetpointChange" message then decides on whether or not to implement requested setpoint change. If it implements the requested change, it sends a "completed" message back to the requesting Agent. Upon receiving both a "requestAction" message and, if appropriate, a "completed" message, the Agent having the "take action" token makes the requested setpoint change and the total rate of operating cost of the entire HVAC system should be reduced.

ADVANCED OPTIMIZATION METHOD (AOM)

In the AOM, each Agent determines the "total rate of operating cost" of the HVAC equipment it is responsible for over the entire range of possible setpoint values. This involves (1) determining how different proposed setpoint changes affect the "rate of operating cost" of the HVAC component it controls, (2) communicating with other Agents to determine how the different proposed setpoint changes affect the "rate of operating cost" of other HVAC components, (3) calculating the expected "total rate of operating cost" for all possible setpoint changes, and (4) finding the proposed setpoint change having the minimum "total rate of operating cost." After each Agent has completed these four tasks, the Agent proposing the setpoint change with the lowest "total rate of operating cost" is allowed by the Coordinator Agent to make that particular setpoint change. This process is then repeated until the Agents involved receive a "no optimization" token from the Coordinator Agent. It is assumed that the optimal operating point is reached when none of the Agents in the optimization process can propose a new setpoint having a lower "total rate of operating cost".

As described above, the AOM requires each Agent to determine the effect of all possible setpoint changes on the "rate of operating cost" for its HVAC and for other HVAC components that may be affected. The latter is accomplished by sending a "inquireCost" messages to other Agents and receiving a "costResponse" messages back for particular setpoint changes.

After an Agent has determined the "total rate of operating cost" for all possible setpoint changes under its control, it is necessary to compare the results with other Agents and then make a decision on which Agent gets to implement the "best" setpoint change. To simplify this comparison, it was decided to use a modified version of the Coordinator Agent that was developed for the SOM. Also, rather than trying to compare "total rate of operating costs" directly, each Agent calculates the decrease in the "total rate of operating cost" between the setpoint having the lowest "total rate of operating cost" and the current setpoint and sends this information, along with the current and proposed new setpoint, to the Coordinator Agent. If there is no setpoint with a lower "total rate of operating cost" than the current setpoint, the Agent sends the current setpoint as the new proposed setpoint to the Coordinator Agent along with a "zero" decrease in the "total rate of operating cost" (i.e., no operating cost savings is possible).

After receiving the current setpoint, proposed new setpoint, and the expected decrease in the total rate of operating cost from all the Agents involved in the optimization, the Coordinator Agent determines which Agent has the largest decrease in total rate of operating cost and sends a message back to that Agent telling it to change the setpoint to the proposed new value. This sequence of Agents determining the effect of all possible setpoint changes, sending a proposed new setpoint corresponding to the largest expected decrease in the total rate of operating cost to the Coordinator Agent, and the Coordinator Agent deciding which setpoint change to implement is repeated as long as the optimization process is underway.

It should be clear from the above discussion that the AOM is considerably more complicated than the SOM. It requires a lot more information to be exchanged between the various Agents and between the Agents and the Coordinating Agent. This in turn makes debugging the testbed program much more difficult. Also, since the testbed does both the HVAC system simulation and the optimization process in the same program, instability problems can occur when changes are made in the various setpoints. For example, if the chiller simulation uses a new chilled water supply temperature setpoint in its chiller performance calculation before the AHU simulations know about the setpoint change, it is possible to get simulated chiller capacities that are larger than the maximum capacity of the chillers. To avoid this problem, the various setpoint changes were implemented in stages so that the different Agents encountered them in the proper sequence. This greatly complicated the implementation of an AOM in the testbed. Also, allowing large setpoint changes can cause system instability problems. A large change in the cooling load can result in a large setpoint change followed by oscillations in the setpoint until the system stabilizes. This was encountered to some degree in the testbed when the internal cooling loads were changed on the hour, even though a time constant on the zone air temperature tended to smooth out these effects. To avoid this problem, the current version of the AOM was constrained to allow setpoint changes of no more than 2.2 ˚C (4.0 ˚F) at any one time. This avoided most of the instability problems, but meant that under certain conditions it took several “optimization steps” to arrive at the “optimal” operating point.

RESULTS

The results obtained using the SOM and the AOM are shown in Table 1 for different Agents doing the optimization. The table presents the cost savings in percent obtained on Day 3 when the Agents performed optimization relative to the case of no optimization. The no optimization case or “reference case” employed fixed values of supply air temperature, chilled water supply temperature, and entering chiller condenser water temperatures of 15.6 ˚C (60.0 ˚F), 5.56 ˚C (42.0 ˚F) and 25.6 ˚C (78.0 ˚F), respectively. Since the savings from the cooling towers optimization was so small (i.e., less than 1%) , it was not combined with the “AHUs + Chiller1” optimization shown below.

Table 1. Cost savings in percent (relative to the reference case with no optimization) for the simplified optimization method (SOM) and the advanced optimization method (AOM) with different Agents doing the optimization.

Simplified Optimization Method - % Savings					
Optimizing Agent(s)	**costAHUFans**	**costChillers**	**costChWPumps**	**costCTFans**	**costTotal**
AHUs	53.64	-0.01	-18.18	-0.01	21.19
Chiller1	-0.10	17.09	-61.99	2.69	5.17
CTs	0.00	-4.37	0.00	31.79	0.79
AHUs + Chiller1	**64.68**	**-11.74**	**-4.45**	**-2.33**	**21.08**

Advanced Optimization Method - % Savings					
Optimizing Agent(s)	**costAHUFans**	**costChillers**	**costChWPumps**	**costCTFans**	**costTotal**
AHUs	53.76	0.08	-18.96	0.08	21.29
Chiller1	0.18	17.08	-62.01	2.66	5.29
CTs	0.19	-4.23	0.03	31.97	0.94
AHUs + Chiller1	**61.71**	**-7.34**	**-8.00**	**-1.01**	**21.72**

It is important to note that an "optimization" performed by the AHU Agents working alone or the Chiller1 Agent working alone is not much of an optimization. Because of the shape of the AHU fan and chiller performance curves, when these Agents consider only the performance of the equipment they are responsible for, the optimal operating point lies at end of the range of allowable setpoints. Thus the AHU Agents push the supply air temperature as low as it is allowed to go and then keep it there. Similarly the Chiller1 Agent pushes the chilled water supply temperature a high as it is allowed to go, where it remains.

A much more interesting optimization is the when the AHU Agents and Chiller1 agents work together to optimize the overall HVAC system performance. As can be seen in Table 1, when the Agents work together they find that the overall cost savings can be maximized by increasing the chiller energy costs (-11.74 % and -7.34 % cost savings for the SOM and AOM methods, respectively) while reducing the AHUs fan energy costs (64.68 % and 61.71 % savings for the SOM and AOM methods, respectively).

For this particular reference case, the combined AHUs + Chiller1 optimization achieves a total cost savings of 21.1 % when the SOM is employed and a total cost savings of 21.7 % when the AOM was used. For the SOM and AOM, respectively, this turns out to be approximately 93 % and 96 % of the maximum possible savings. The maximum possible savings of approximately 22.7 % was found by running the testbed for fifteen chilled water supply temperature setpoints between 2.22 ˚C (36.0 ˚F) and 11.1 ˚C (52.0 ˚F) (with the supply air temperature setpoint set 5.56 ˚C (10.0 ˚F) above the chilled water supply setpoint) and at each time step selecting the lowest total operating cost. The lowest operating costs at each time step were then summed to determine the minimum possible (or optimal) operating cost.

The total rate of operating cost and the rate of operating cost of the different HVAC components for the reference case (no optimization) is shown for Day 3 on the left hand side of Figure 6. The right hand side of Figure 6 shows the same costs for the AOM method when the AHU and Chiller1 Agents are acting together to do the system optimization. Days 1 and 2, which are not shown, were used for performance identification.

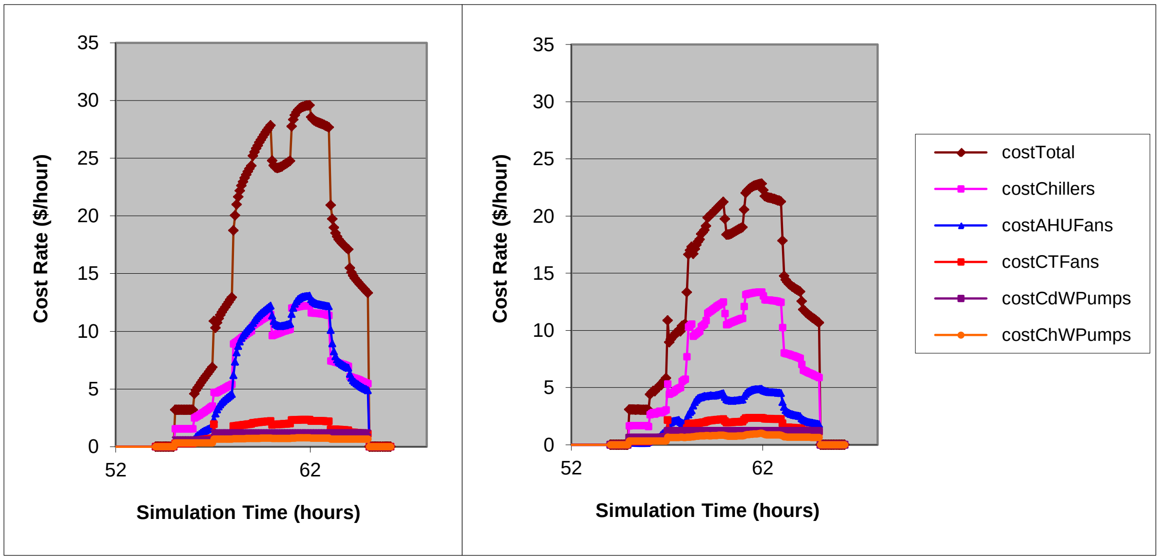

Figure 6. "Rate Of Operating Costs" for the reference case (no optimization) is shown for Day 3 on the left side of this figure. The same costs for the AOM method with the AHU and Chiller1 Agents working together to do the system optimization is shown on the right hand side.

Figure 7 shows the resulting changes in the chilled water supply temperature setpoints for the third simulation day (Day #3) when the AHU1, AHU2, and Chiller1 Agents worked together to optimize the overall system performance using the SOM and AOM. Also shown in this Figure is the "optimal" chilled water supply setpoint found from running the fifteen cases described above and at each time step selecting the chilled water supply setpoint having the lowest rate of total operating cost. From approximately 11:00 AM to approximately 5:00 PM, the optimal chilled water supply temperature setpoint was found to be 4.44 ˚C (40.0 ˚F). During this same period, the AOM found the optimal setpoint to be 3.89 ˚C (39.0 ˚F), while the SOM found it to be 3.33 ˚C (38.0 ˚F). These differences from the true optimal setpoint that resulted from employing the AOM and SOM are likely to be caused by slight inaccuracies in identified chiller performance or the identified supply air fan performance and/or, in the SOM case, by the fact that setpoint changes were made in 1.1 ˚C (2.0 ˚F) increments.

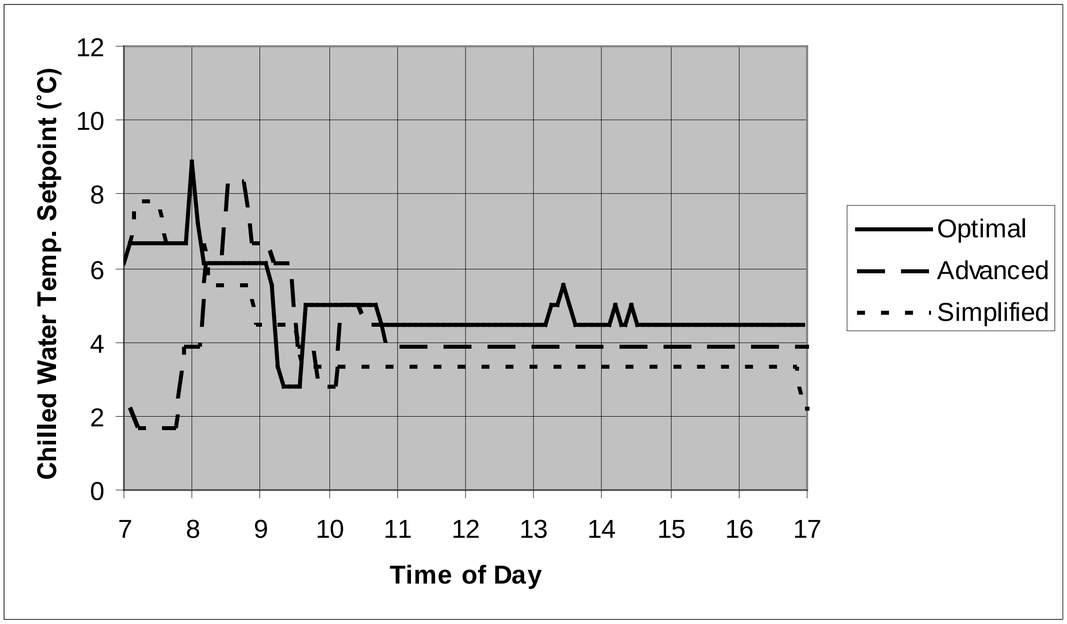

Figure 7. True optimal chilled water supply temperature setpoints (solid line) for simulation Day #3 and the "near optimal" setpoints found by the Chiller1 and AHU Agents working together using the SOM and the AOM.

While the total cost savings of 21.1% and 21.7% achieved by AHU and Chiller1 Agents working together using the SOM and AOM, respectively, appear to be significant, they are strongly dependent on the reference case chosen and cannot be considered representative of the actual savings that might be achieved in a real building/HVAC system. For example, if the reference case was changed so that the supply air temperature setpoint was 12.8 ˚C (55.0 ˚F) instead of 15.6 ˚C (60.0 ˚F) and the chilled water supply temperature and the entering chiller condenser water temperatures setpoints were kept, respectively, at 5.56 ˚C (42.0 ˚F) and 25.6 ˚C (78.0 ˚F), the total cost savings achieved by the AHUs and Chiller1 Agents working together using the AOM would only be 6.3%. This is because the maximum (true optimal) total cost savings possible with this reference case is only 7.9%.

Conclusions and Recommendations

Although the limitations of the current testbed (e.g., a single simulation day, a fixed ambient air wetbulb temperature, a simplified approach to handling cooling coil performance, etc.) make it difficult to draw definitive conclusions regarding the potential energy savings resulting from the use of intelligent agents, the current work clearly demonstrates that the concept of using intelligent agents to optimize the performance of a real building/HVAC system is extremely promising. The testbed showed that piecewise linear regression appears to work well for identifying the performance of HVAC equipment and that intelligent agents can work together to achieve "near optimal" results. The use of intelligent agents make it possible to solve the problem of building system optimization in a "distributed manner" and avoids many of the problems associated with optimization methods that require that information on the performance and status of all HVAC systems and equipment be available in one location (i.e., at the device doing the optimization).

Both the simplified optimization method (SOM) and the advanced optimization method (AOM) discussed in this paper performed almost equally well in reducing the operating cost of the simulated building HVAC systems. For both methods, when the AHU and Chiller1 Agents acted together to do the optimization, the total "one day" cost savings was approximately 21.1% and 21.7%, respectively, for the reference case (no optimization case) having a supply air temperature, chilled water supply temperature, and entering chiller condenser water temperature setpoints of 15.6 ˚C (60.0 ˚F), 5.56 ˚C (42.0 ˚F) and 25.6 ˚C (78.0 ˚F), respectively. For the SOM and AOM, respectively, this was approximately 93% and 96% of the maximum possible savings. However, since these cost savings are strongly dependent on the reference case chosen, they cannot be considered representative of the actual savings that might be achieved in a real building/HVAC system.

As discussed above in the section on the Advanced Optimization Method, the implementation of the AOM was not easy. The AOM involves the exchange of large amounts of information among Agents and can result in very large setpoints changes. The latter can cause instability problems when disturbances, such as cooling load changes, occur. The SOM on the other hand, is considerably easer to implement. It involves Agents taking turns in optimizing the particular piece of HVAC equipment they are responsible for and requires only small changes in the setpoints. For the optimization cases discussed in this report, both methods achieved almost the same cost savings. Because of this, it is strongly recommended that, in any future studies involving the use of intelligent agents to optimize building/HVAC performance, consideration first be given to implementing the SOM rather than the AOM.

Finally, since simulation studies can never fully replace the need for "real world" experience, the work and conclusions of this paper need to be verified by future studies under controlled laboratory conditions and in real building/HVAC systems.

REFERENCES

ASHRAE. 2007. 2007 ASHRAE Handbook—HVAC Applications, Chapter 41. Atlanta: American Society of heating, Refrigerating and Air-Conditioning Engineers, Inc.

Barr, A, Fleigenbaum E. A., 1981). The Handbook of Artificial Intelligence, Volume 1, HeurisTech Press, Stanford CA

Barr, A, Fleigenbaum E. A., 1982). *The Handbook of Artificial Intelligence, Volume 2*, HeurisTech Press, Stanford CA

Braun, J. E., Klein, S. A., et al. (1989a). Methodologies for Optimal Control of Chilled Water Systems Without Storage. ASHRAE Trans. 95: 1.

Braun, J. E., Klein, S. A., et al. (1989b). Applications of Optimal Control to Chilled Water Systems Without Storage. ASHRAE Trans. 95: 1.

Bushby, S. T., (1997). BACnet™ - A Standard Communication Infrastructure for Intelligent Buildings. *Automation in Construction*, Vol. 6 No. 5-6, p. 529-540.

Cumali, Z., (1988). Global Optimization of HVAC System Operations in Real Time. ASHRAE Trans. 94: 1.

House, J. M., Smith. T. F., et al. (1991). Optimal Control of a Thermal System. ASHRAE Trans. 97: 2.

Ljung, L., (1987) *System Identification Theory for The User*, Prentice-Hall, Inc., Elglewood Cliffs, NJ

Pantic M. et al, (2005). Teaching Introductory Artificial Intelligence Using a Simple Agent Framework. IEEE Trans. On Education. Vol. 48

Pape, F. L. F., Mitchell, J. W., et al. (1990). Optimal Control and Fault Detection in Heating, Ventilating, and Air-Conditioning Systems. ASHRAE Transactions 97: 1.

Wang, Y., Cai, W., et al. (2004). A simplified modeling of cooling coils for control and optimization of HVAC systems. Energy Conversion and Management 45, p. 2915-2930.

Appendix A

As described in the section entitled Identification, both the SOM and the AOM allowed for a simplifying approach to be employed in carrying out joint chiller-AHU optimization. It involved not permitting the difference between the supply air temperature setpoint and the chilled water temperature setpoint to be less than 5.56 ˚C (10.0 ˚F). This simplification is based upon a separate analysis of a typical cooling coil, which showed that (1) for a properly sized coil and large cooling loads, a supply air temperature set point 5.56 ˚C (10.0 ˚F) above the chilled water temperature setpoint was either optimal or very close to optimal, and (2) for moderate and low cooling loads, while the optimal supply air temperature setpoint was closer to the chilled water temperature setpoint, the effect on the cost of operation of keeping it 5.56 ˚C (10.0 ˚F) above the chilled water temperature was very small.

Since the very simple cooling coil model employed in the testbed did not allow for an adequate check on the validity of this simplified approach, a separate analysis was conducted using a simplified cooling coil model developed by Yao-Wen Wang (Wang 2004) for "control and optimization of HVAC systems". Wang derived a heat transfer equation for a HVAC cooling coil given by:

$$Q := \left[\frac{C1 \cdot (\text{rateSA})^{e}}{1 + C2 \cdot \left(\frac{\text{rateSA}}{\text{rateChw}} \right)^{e}} \right] \cdot (\text{tempAirIn} - \text{tempChw})$$

where:

Q = cooling coil load
rateSA = supply air mass flow rate,
rateChw = chilled water mass flow rate,
tempChw = temperature of chilled water entering the cooling coil, and
tempAirIn = entering air wet bulb temperature for a wet coil.

The coefficients C1, C2 and e were set equal to 2.439, 0.498, and 0.8, respectively. These values were determined experimentally by Wang, for a cooling coil with a cooling capacity reasonably close to that of the cooling coil used in the testbed (400 kW vs. 470 kW, respectively).

Also, since the cooling coil model used in the testbed does not consider latent cooling loads, it was assumed that the cooling coil model developed by Wang was applicable to dry cooling coils and the variable tempAirIn in the above equation was set equal to the entering air dry bulb temperature.

After defining a part load factor, plf , given by:

plf = Q/Qmax,

where Qmax is the cooling capacity of the cooling coil used in the Testbed, and using the fact that

rateSA = plf*Qmax / (cp*(tempAirIn – tempSA)),

where cp is the specific heat capacity of air,

and setting tempAirIn equal to 24.0 ˚ (75.2 ˚F),

the above equations for Q and rateSA were solved to give the chilled water flow rate, rateChw, for different values of plf, tempChw, and tempSA.

Figure A1 shows the calculated values of rateChw for plf = 1, three values of tempChw (3.33 ˚C (38.0 ˚F), 5.56 ˚C (42.0 ˚F), and 8.33 ˚C (47.0 ˚F)) and for values (on the x-axis) of tempSA between 1.0 ˚C (1.8 ˚F) and 10 ˚C (18 ˚F) above tempChw. Note that at this plf the rateChw rises rapidly as tempSA approaches tempChw.

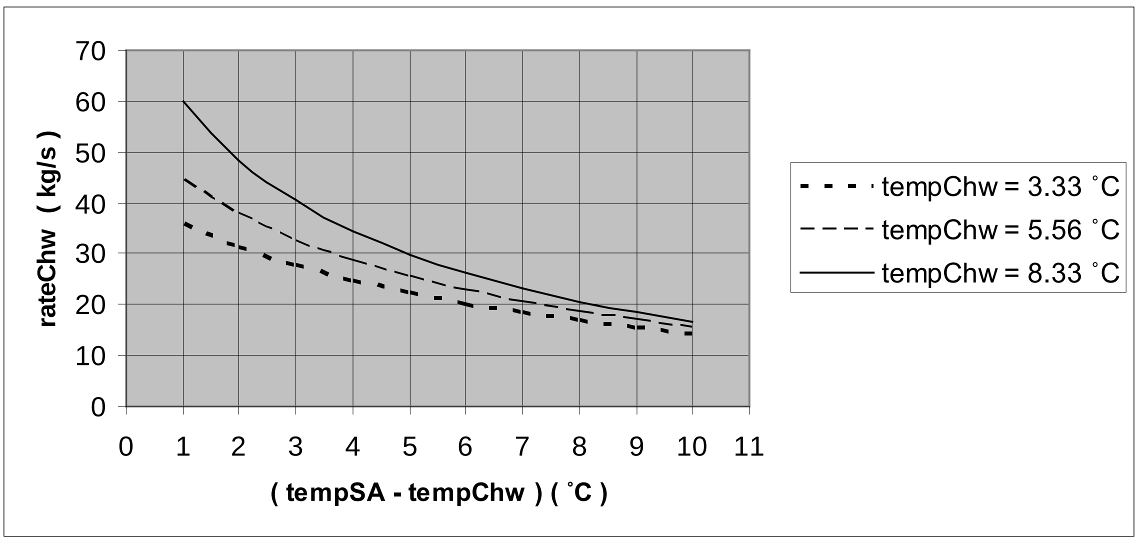

Figure A1. Chilled water mass flow rate for plf = 1 and tempChw equal to 3.33 ˚C (38.0 ˚F), 5.56 ˚C (42.0 ˚F), and 8.33 ˚C (47.0 ˚F).

Employing the fan and pump performance curves used in the testbed and assuming an electricity cost of $0.09/kWh, the combined rate of operating cost of an AHU supply air fan and a chilled water pump can be calculated. (For moderate to large cooling loads, there will be two cooling fans and two chilled water pumps running. Also, the operating cost of the chillers can be ignored because it is constant for specific values of plf and tempChw).

Figures A2, A3, and A4 show the rate of operating cost of operating an AHU supply air fan and a chilled water pump for plf equal to 1.0, 0.7, and 0.3, respectively. Each figure shows the results for three values of tempChw (3.33 ˚C (38.0 ˚F), 5.56 ˚C (42.0 ˚F), and 8.33 ˚C (47.0 ˚F)) and for values (on the x-axis) of tempSA between 1.0 ˚C (1.8 ˚F) and 10 ˚C (18 ˚F) above tempChw.

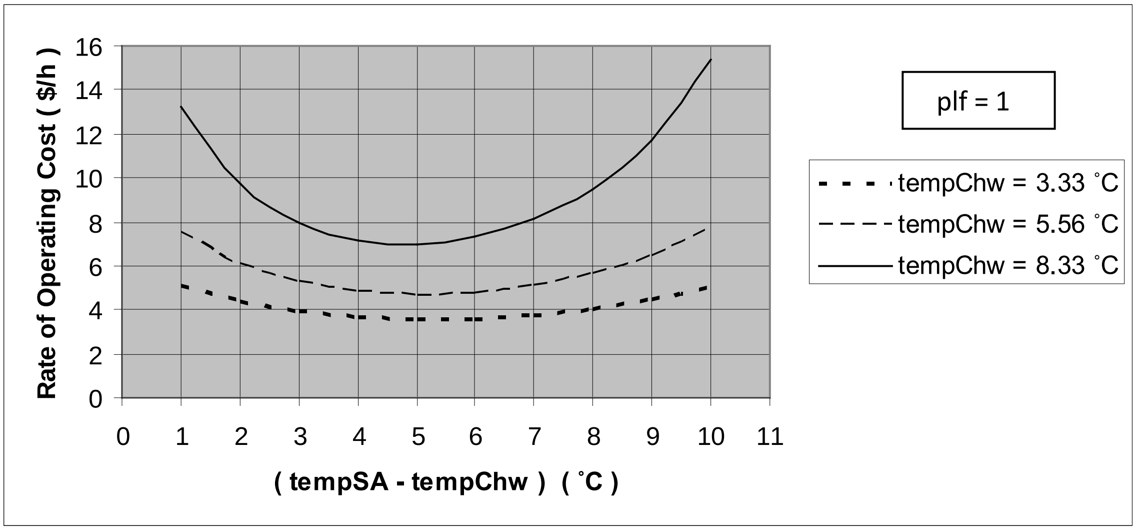

Figure A2. Rate of Operating Cost for a AHU Supply Air Fan and a Chilled Water Pump for plf = 1 and tempChw equal to 3.33 ˚C (38.0 ˚F), 5.56 ˚C (42.0 ˚F), and 8.33 ˚C (47.0 ˚F).

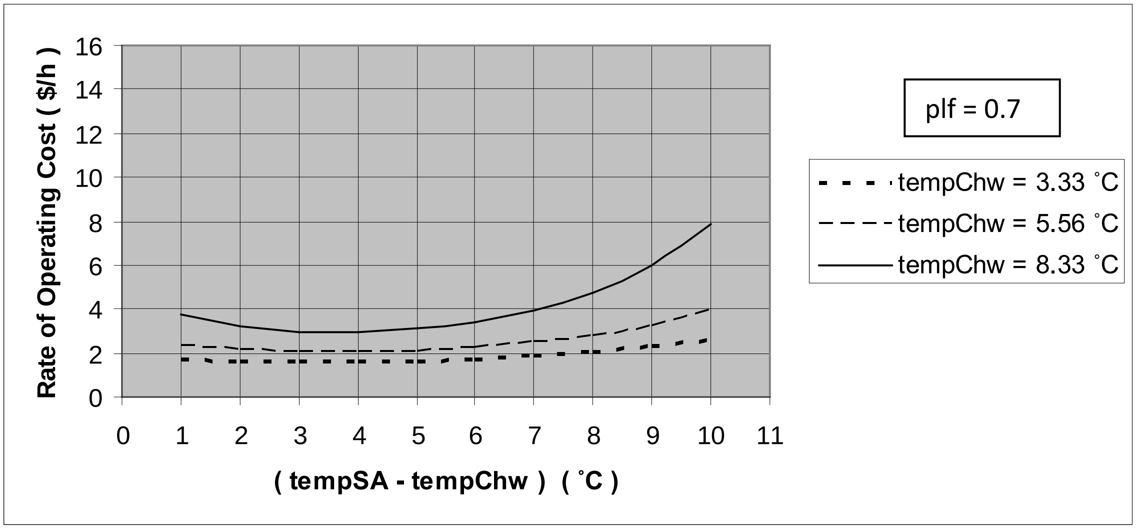

Figure A3. Rate of Operating Cost for a AHU Supply Air Fan and a Chilled Water Pump for plf = 0.7 and tempChw equal to 3.33 ˚C (38.0 ˚F), 5.56 ˚C (42.0 ˚F), and 8.33 ˚C (47.0 ˚F).

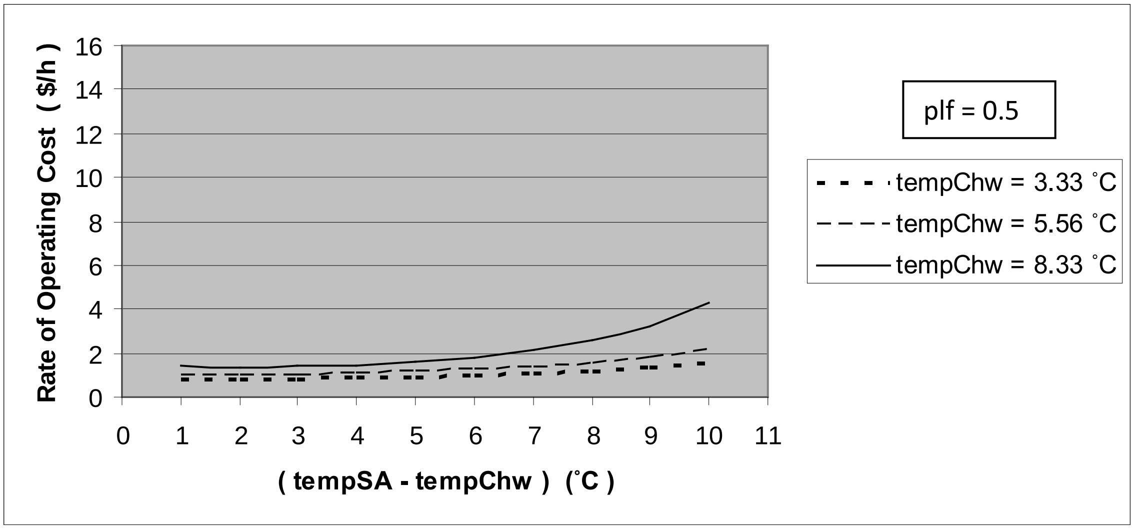

Figure A4. Rate of Operating Cost for a AHU Supply Air Fan and a Chilled Water Pump for plf = 0.5 and tempChw equal to 3.33 ˚C (38.0 ˚F), 5.56 ˚C (42.0 ˚F), and 8.33 ˚C (47.0 ˚F).

As shown in Figure A2, for a plf = 1.0, the minimum cost for the thee different tempChw shown occurs when tempSA is between 4.4 ˚C (8.0 ˚F) and 5.56 ˚C (10.0 F) above tempChw. However, the minimum is very flat and there is little difference in the operating cost within this range.

For plf equal to 0.7 and 0.5, Figures A3 and A4 show that the minimum is not very well defined for moderate and small cooling loads. These figures also show that for these loads, choosing a tempSA 5.56 ˚C (10.0 ˚F) above tempChw does not have a significant impact on the cost of operation.

These results substantiate the assumption used in developing the Testbed that an optimization method (either SOM or AOM) that does not permit the supply air temperature setpoint to be less than 5.56 ˚C (10.0 ˚F) above the chilled water sepoint will give near optimal results.