A Multi-Agent Control Approach for Optimization of Central Cooling Plants

Rita C. Jaramillo
Purdue University, United States of America, rjaramil@purdue.edu

James E. Braun
Purdue University, United States of America, jbraun@purdue.edu

W. Travis Horton
Purdue University, United States of America, wthorton@purdue.edu

A Multi-Agent Control Approach for Optimization of Central Cooling Plants

Rita JARAMILLO[1*], James BRAUN[2], Travis HORTON[3]

[1]*School of Mechanical Engineering, Purdue University, West Lafayette, IN
E-mail: rjaramil@purdue.edu

[2]Professor of Engineering, School of Mechanical Engineering
Purdue University, West Lafayette, IN
E-mail: jbraun@purdue.edu

[3] Assistant Professor of Civil Engineering with Courtesy Appointment in Mechanical Engineering
Purdue University, West Lafayette, IN
E-mail: wthorton@purdue.edu

* Corresponding Author

ABSTRACT

This paper presents an application of a multi-agent control approach for supervisory control of large central cooling plants. The starting point for this work was a multi-agent control simulation framework developed by Cai (2015). To adapt the framework to this problem, agents representing the performance of the different devices of the plant were developed and an optimization method capable of handling non-convex functions and discontinuous design spaces was developed and incorporated in the framework. A case study of an existing cooling plant was utilized to evaluate the approach in terms of optimality and computational resources. Simulations were carried out for different performance conditions to predict the performance of the plant under three different control strategies: 1) multi-agent control, 2) centralized optimization based on mathematical programming techniques and 3) a heuristic control strategy. The results showed that significant savings can be achieved through the implementation of multi-agent control. It is expected that, if each hardware component of the plant comes with an integrated agent that represents its behavior, then the proposed multi-agent framework could automatically generate the multi-agent structure and control algorithm after some relatively simple pre-configuration steps. This will reduce the site-specific engineering and will provide a more economic and easy to configure solution for central cooling systems.

INTRODUCTION

A large central cooling system consists of several chillers, cooling towers and pumps that supply chilled water to satisfy the cooling requirements of one or more buildings. Optimal supervisory control of such systems involves the determination of the mode of operation and set points that minimize operating costs while satisfying cooling and comfort requirements. The problem is complicated because of the presence of both continuous and discrete control variables. Most of the research related to optimal supervisory control of central cooling systems that has been conducted in the last three decades has focused on centralized control approaches, such as Braun et al. (1989a, b), Ahn and Mitchell (2001); Yao et al. (2004); Torzhkov et al. (2010), and Zhang and Turner (2012). Although these studies have demonstrated the effectiveness of optimal or near-optimal control in reducing operational costs, the results have not been widely implemented. Some limitations of centralized optimal control are the need for detailed information on the performance profiles of the cooling plant equipment in order to build a model for the optimization process, and, once implemented, the plant model and control sequences will need to be updated by experts every time a modification (such as the introduction of new equipment) is made to the plant.

A promising approach that addresses limitations of earlier approaches is the implementation of distributed multi-agent-based optimal control. The use of intelligent agents makes it possible to solve the optimization problem in a distributed manner by breaking a big complex problem into smaller, more manageable pieces that can be solved independently and in parallel by individual agents. The individual solutions can then be handled by a coordination agent that achieves some consensus. Since intelligent agents can solve individual problems to optimize performance without having total knowledge of the system, they would also add adaptive capability to the control system, i.e., the system could be more easily reconfigured to adapt to changes such as the introduction of new equipment. However, some drawbacks of this approach are the additional data transfer equipment required and the optimality traded off for reduced computations. Many approaches for distributed control and multi-agent systems have been proposed and demonstrated in a number of fields. However, a review of multi-agent control shows relatively few applications in the HVAC field (Treado, 2010; Sun et al, 2010; Kelly and Bushby, 2012). In these studies, results that document the performance of the control algorithms are very scarce and the proposed multi-agent control strategies were validated using simulations on small-scale HVAC systems with very simple models on the cooling plant side.

This paper presents an application of a multi-agent control approach for supervisory control of large central cooling plants. The work starts from a multi-agent control simulation framework developed by Cai (2015) for optimization-based supervisory control of distributed air-conditioning systems. In this setting, assuming that each hardware component of a system has an integrated agent that represents its behavior, then the framework can automatically generate the multi-agent structure and control algorithm after some relatively simple pre-configuration steps, reducing site-specific engineering. Although the proposed framework provides good flexibility in design of control topology it has some limitations and it might not be directly applied to some kinds of equipment: the distributed consensus-based algorithms utilized are conceived for convex functions and continuous design spaces. Therefore, they are not good at handling discrete variables such as multiple operating modes that are often present in HVAC systems. To adapt the framework to the current problem, agents representing the performance of the different devices of central cooling plants were developed and inserted in the framework and an optimization method capable of handling non-convex functions and discontinuous design spaces (Genetic Algorithm) was developed and incorporated in the framework. A case study of a cooling plant without significant storage was addressed to conduct a simulated demonstration of the approach for different operating conditions. The results in terms of optimality and computational resources were compared with two benchmarks: centralized optimization with conventional mathematical programming techniques and a heuristic control strategy.

1. MULTI-AGENT APPROACH

The multi-agent control approach utilized here consists of two main elements: a multi-agent simulation framework, which includes component agents representing the behavior of each physical component of the plant, and two optimization algorithms: a consensus-based distributed algorithm originally included in the framework called the alternating direction method of multipliers (ADMM) and a genetic algorithm (GA). The GA was developed and incorporated in the framework to provide an alternative for finding the global optimal operating point of a system in the presence of non-convex functions and discontinuous design spaces.

1.1. Component Agents

Previous work developed by Jaramillo et al. (2014) presented the development of a mathematical model of a cooling plant using MATLAB software. In this setting, each hardware component of the plant (i.e. chillers, cooling towers and pumps) was represented using semi-empirical models as a separate set of mathematical relationships with its own parameters, inputs and output variables. The parameters of the models were determined from regression of performance data. From there, agents representing each component behavior can be easily created. Each agent consists of a cost function (power consumption in this case) and a set of constraints. Once created, the agents can be incorporated into the multi-agent framework and interconnected according to the arrangement of the physical plant.

1.2. Multi-agent Framework

The multi-agent control framework developed by Cai et al. (2015) was built in MATLAB as a simulation prototype that can be replicated in other programming environments to support hardware implementation. The framework defines a general component agent structure as well as the flow connections between agents. A general component agent is written as a super class from which each component class can inherit the basic agent structure. The properties of the agent class consist of a collection of cost functions, and linear and non-linear equality and inequality constraints that characterize the behavior of a specific hardware component. The cost functions could be actual power consumption that needs to be minimized or some other performance metrics. Another important property of each component agent is the agent's group number. This parameter is used in the setup of the distributed optimization-based controllers: all the component agents with the same group number will be assigned to one local optimizer controller. The grouping might depend on physical distances among the devices, function, network structure or other criteria. A centralized controller will be synthesized if all the component agents are assigned the same group number. In this setting, the procedure to create a multi-agent control system is straight forward: assuming that all the component agents are available, one can simply drag and drop them in a project canvas and then specify the inter-agent connections, which are stream variables. Figure 1 illustrates the procedure to create a multi-agent control system for a central cooling plant.

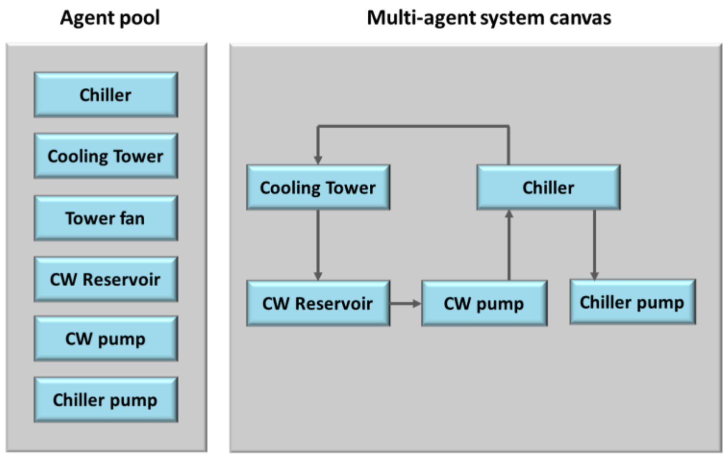

Figure 1. Procedure to Set-up a Multi-agent System. Adapted from Cai et al. (2015)

Once the component agents and their connections are defined, the framework will automatically compile the code and compose the optimization problem according to the specified configuration. The compilation process consists of several steps such as extraction of the cost functions of the different agents to construct the total cost function and the elimination of redundant equality constraints. These steps will be carried out for each group of components. Then, for each group *i* the composed optimization problem takes the form of Eq. 1.

$$\begin{aligned} & min_{X_i} f_i(X_i) \qquad i = 1, \dots, N \\ & subject\ to \quad A_i X_i \leq B_i \\ & \qquad A_{eq,i} X_i = B_{eq,i} \\ & \qquad g_i(X_i) \leq 0 \\ & \qquad h_i(X_i) = 0 \end{aligned} \tag{1}$$

In the above expression X_i is a vector of the local variables of sub-problem *i* and *N* is the number of subproblems. To complete the process of composing the distributed problem, consensus constraints need to be specified. These constraints enforce the local copies of the same variable to match between different agents. Taking the system in Figure 1 as an example, the water flow leaving the condenser pump is the same entering the chiller's condenser. If these two devices are assigned different group numbers, there will be two local variables corresponding to the same water flow properties. The minimization of the variable-speed pump power will favor lower condenser water flow while the minimization of the chiller power will be benefited by higher condenser water flow. So the minimization of power consumption of each group will drive two local variables representing the same physical quantity in opposite directions. Therefore, consensus constraints of the form expressed in Eq. 2 are necessary to enforce these local variables to converge to the same value.

$$X_i = F_i Z \qquad i = 1, \dots, N \tag{2}$$

where Z is a vector that contains the global variables of the problem and F_i is a matrix that picks out the variables of Z that correspond to X_i.

1.3. Multi-agent Optimization Algorithms

Distributed Optimization: Alternating Direction Method of Multipliers

The ADMM is an augmented Lagrangian method for solving distributed consensus problems that was introduced in the 1970s. A description of the method can be found in Summers and Lygeros (2012). The method was adapted to solve problems of the form expressed in Eq. 1 and Eq. 2. The augmented Lagrangian for these kind of problems is given in Eq. 3

$$L = \sum_{i=1}^{N} L_i(X_i, Z, Y_i) = \sum_{i=1}^{N}\Big(f_i(x_i) + Y_i^T(X_i - F_i Z)\Big) + \frac{\sigma}{2}\|X_i - F_i Z\|_2^2 \tag{3}$$

Where Y_i are vectors of the Lagrange multipliers and σ is a penalty parameter. The ADMM algorithm consists of the iterations

$$\begin{aligned} X_i^{k+1} &= argmin_{X_i} L_i\big(X_i^k, Z^k, Y_i^k\big) \quad s.t. \quad X_i \in \boldsymbol{C}_i \\ Z^{k+1} &= argmin_Z L_i\big(X_i^{k+1}, Z^k, Y_i^k\big) \\ Y_i^{k+1} &= Y_i^k + \sigma\big(X_i^{k+1} - F_i Z^{k+1}\big) \end{aligned} \tag{4}$$

The method alternatively minimizes X and Z, which allows the X_i minimizations to be done in parallel. In this particular form of the problem, the minimization of each variable of the vector Z reduces to an averaging of the equivalent local variables and Lagrange multipliers as expressed in Eq. 5.

$$(Z^{k+1})_j = \frac{1}{N_j}\sum_{i=1}^{N}\left(\big(X_i^{k+1}\big)_j + \frac{\big(Y_i^k\big)_j}{\sigma}\right) \tag{5}$$

With this distributed formulation, the original optimization problem is fragmented into several sub-problems with reduced dimensions and less constraints, which can be solved in parallel. A hardware implementation of this distributed decision making process is shown in Figure 2. The bottom layer corresponds to the sensing network that collects the required operating conditions. Above the sensing layer is a component agent layer that includes the agents representing the behavior of all devices. On top of the component agent layer, there is an optimizer agent layer, which is responsible for solving each sub-problem. Each optimizer agent calls the related component agents iteratively to optimize its corresponding cost function independently and in parallel with the other optimization agents. The consensus requirements among the local variables are enforced by a coordination layer that collects the local copies of all the variables, updates the dual variables accordingly and feeds them back to the optimizer agents to let them re-optimize with respect to the updated information. The iteration process continues until the termination criteria are met.

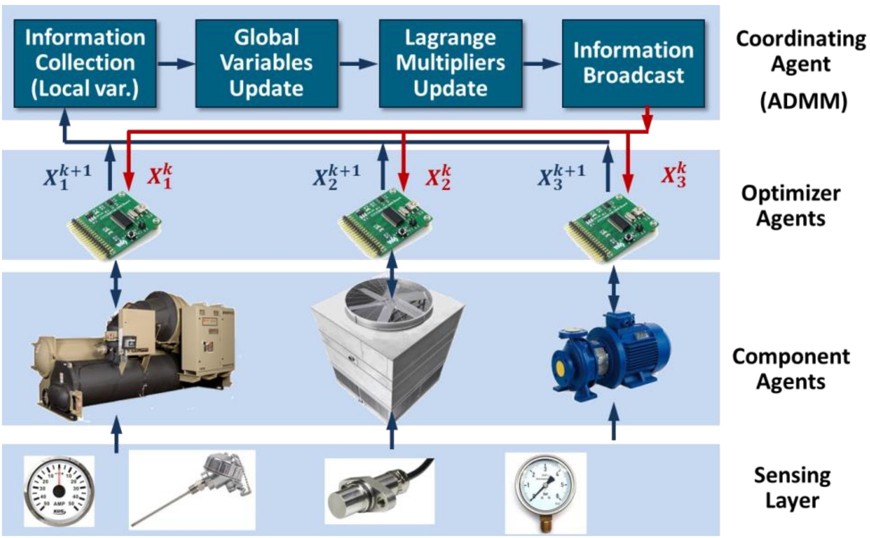

Figure 2. Hardware Implementation of the Multi-agent Controller with ADMM

Centralized, parallel optimization with Genetic Algorithm

Genetic algorithms (GA) are stochastic methods based on the principles of genetics and natural selection. A generic GA starts with a population of N_p candidate solutions randomly generated from the search space. An objective function is used to quantify the fitness of each individual and genetic operators of selection, crossover and mutation are applied to produce a new generation of individuals. The process of populating new generations continues until certain convergence criteria are satisfied. The algorithms for the genetic operators depend on the encoding method and the application. The GA conceived for optimization in the context of the multi-agent framework is real coded and has incorporated certain characteristics to make it suitable for solving large dimension problems in a more efficient way. Some of those features, based on the work presented by Zhu et al. (2014) are the generation of the initial population as a uniform array of points that covers the search space delimited by bounds and linear constraints and an elitist strategy that combines all the individuals from the previous generation and the new population, and selects the fittest individuals to constitute the new generation. One disadvantage of GAs is that their performance is not always satisfactory in the presence of equality constraints. To sort this difficulty, the proposed GA was combined with Broyden's method (a quasi-Newton method) to handle the non-linear equality constraints that arise from the behavior of the components.

2. CASE STUDY

The case-study considered in this work is a simplified version of the Northwest Chiller Plant at the main campus of Purdue University in West Lafayette, IN. The plant delivers chilled water through 37 Km of underground piping to partially meet the cooling requirements of more than 150 buildings on the campus (an average of 92.4 MMTon-hr per year). A schematic of the simplified plant, illustrating all the variables of interest is shown in Figure 3. The simplified case-study consists of three centrifugal chillers, each one with 2000 Ton nominal cooling capacity to produce chilled water for the campus, an evaporative counter-flow cooling tower with three cells, and three variable speed condenser water pump each one with 6000 gpm nominal flow rate. In the chilled water loop, the water returning from campus is cooled as it is circulated through the chiller and then returned to campus by high-pressure pumps. The chilled water pumps are considered part of the campus chilled water distribution systems and do not add to the costs of operation of the plant, therefore they will not be included in the model. In the condenser water loop, the chillers reject heat to the water that is circulated through the cooling tower cells and stored in the cold well, a common reservoir with capacity of 90,000 gal (341m^3). From there, the water is pumped again to the chillers by the condenser water pumps.

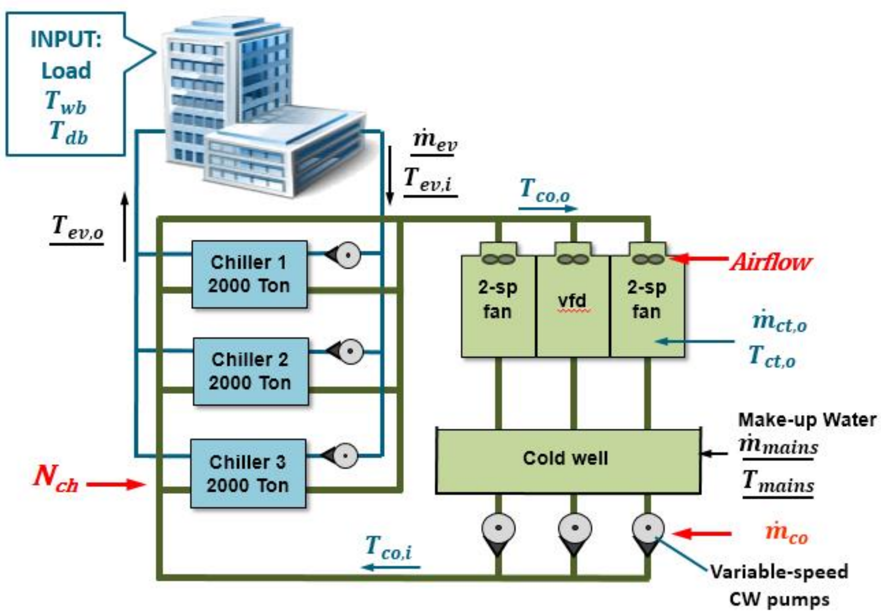

Figure 3. Schematic of the cooling plant

The baseline control for the plant involves the following heuristic strategies: (1) The chillers are sequenced based on the operators experience to meet the cooling load. (2) The condenser water pumps are sequenced with the chillers so that each chiller's condenser operates approximately with the nominal design flow. (3) There is feedback control of the tower fans to maintain a constant condenser water supply temperature set-point of 73°F, even though the minimum inlet water temperature for each chiller condenser is 55°F.

Optimal supervisory control of such a plant involves determining the values of the control variables that minimize the total power consumption at any time in response to uncontrolled variables (ambient dry-bulb and wet-bulb temperatures, and cooling load). Three independent control variables are considered in this problem: the total cooling tower air flow (*Airflow*), the total condenser water flow ($\dot{m}_{co}$) and the number of active chillers (N_{ch}) at any time. Other dependent optimization variables (in blue) appear as a result of the distributed formulation of the problem. The boundary conditions are the underlined variables in Figure 4. The chilled water supply set-point ($T_{ev,o}$) is fixed at 43°F, and the chilled water supply flow rate ($\dot{m}_{ev}$) which is fixed according to the nominal chiller's evaporator flow rate (3200 gpm per active chiller).

2.1. Optimization problem composition

Distributed formulation

If the component agents are grouped according to its kind, four groups are formed and a distributed formulation is synthetized by the framework as shown in Eq. 6 to Eq. 10. The functions in the expressions correspond to the cost functions (power consumption) and constraints related to the behavior of each device of the plant.

Group 1: Chillers

$$min_{\left[\dot{m}_{co}^{(1)},T_{co,i}^{(1)},T_{co,o}^{(1)}\right]}\left\{Chiller_{pow}(\dot{Q}_{ev},\ T_{ev,o},T_{co,o}\ )\right\} \tag{6}$$

$$\text{Subject to:}\quad T_{co,o}^{(1)} = Chiller_{Tco,o}\left(\underline{\dot{Q}_{ev}},\ \underline{T_{ev,o}},\dot{m}_{co}^{(1)},T_{co,i}^{(1)}\right)$$

Group 2: Cooling tower cells

$$min_{\left[Airflow^{(2)},\dot{m}_{co}^{(2)},T_{co,o}^{(2)},\dot{m}_{ct,o}^{(2)},T_{ct,o}^{(2)}\right]}\left\{Power_{ct}\right\} \tag{7}$$

$$\text{Subject to:}\quad T_{ct,o}^{(2)} = CoolingTower_{T_{ct,o}}\left(\underline{T_{wb}},\underline{T_{db}},Airflow^{(2)},\dot{m}_{co}^{(2)},\ T_{co,o}^{(2)}\right)$$

$$\dot{m}_{ct,o}^{(2)} = CoolingTower_{\dot{m}_{ct,o}}\left(\underline{T_{wb}},\underline{T_{db}},Airflow^{(2)},\dot{m}_{co}^{(2)},T_{co,o}^{(2)}\right)$$

Group 3: Condenser water pumps

$$min_{\left[\dot{m}_{co}^{(3)}\right]}\left\{Power_{pump}\right\} \tag{8}$$

Group 4: Cold well

$$min_{\left[\dot{m}_{co}^{(4)},T_{co,i}^{(4)},\dot{m}_{ct,o}^{(4)},T_{ct,o}^{(4)}\right]}\{0\} \tag{9}$$

$$\text{Subject to:}\quad T_{co,i}^{(4)} = Coldwell\left(\dot{m}_{co}^{(4)},\ \dot{m}_{ct,o}^{(4)},T_{ct,o}^{(4)},\underline{T_{mains}}\right)$$

The consensus constraints from this formulation are:

$$\begin{gathered} Airflow = Airflow^{(2)} \\ \dot{m}_{co} = \dot{m}_{co}^{(1)} = \dot{m}_{co}^{(2)} = \dot{m}_{co}^{(3)} = \dot{m}_{co}^{(4)} \\ T_{co,i} = T_{co,i}^{(1)} = T_{co,i}^{(4)} \\ T_{co,o} = T_{co,o}^{(1)} = T_{co,o}^{(2)} \\ \dot{m}_{ct,o} = \dot{m}_{ct,o}^{(2)} = \dot{m}_{ct,o}^{(4)} \\ T_{ct,o} = T_{ct,o}^{(2)} = T_{ct,o}^{(4)} \end{gathered} \tag{10}$$

In this distributed formulation, the original centralized problem is fragmented into four sub-problems with reduced dimensions and less constraints, which can be solved in parallel.

Genetic Algorithm Formulation

The GA is intended to be compatible with the multi-agent framework topology; therefore, the distributed formulation expressed in Eq. 6 to Eq. 9 can be also utilized to solve the problem with the GA. In this context, the cost function is the sum of the cost functions of each optimizer agent; therefore, the optimization problem is solved in a centralized manner by a coordinator agent, while the computation of the objective function and non-linear constraints is distributed and executed in parallel by the different optimizer agents.

3. OPTIMIZATION RESULTS

This section presents the results of the multi-agent approach applied to the case study described above. Four different operating conditions of the plant were considered as shown in Table 1. The optimization results are compared with two benchmarks: the heuristic control strategy previously described and centralized optimization with mathematical programming techniques. It is important to note that the problem considered here involves a discontinuous control variable (number of active chillers) and the cost function (power) is non-convex. In such cases, the convergence of the solution to the optimum point cannot be guaranteed. Therefore, performance maps of the plant were elaborated for each of the operating conditions considered to obtain the "true optimum" operating point. This optimum point was used as a baseline to assess the convergence of the different optimization methods. Table 1 presents the total power consumption obtained with the heuristic control strategy, and the power consumption and savings corresponding to the "true optimum" operating point. The savings were evaluated with respect to the heuristic strategy.

Table 1. Optimum operating point of the Cooling Plant and power savings for four Operating Conditions.

Operating conditions			Heuristics	"True optimum" operating point			
Test Nr	Load Ton	Wet bulb °F	Power kW	Active Chillers	Power kW	Savings kW	Savings %
1	2000	50	1145.46	2	1028.88	116.58	11.33
2	2000	80	1351.06	1	1302.61	48.45	3.72
3	4000	50	2324.22	3	1989.87	334.35	16.80
4	4000	80	2702.12	2	2612.62	89.50	3.43

The optimization problem was solved using three methods: 1) Centralized optimization using mathematical programming techniques (MATLAB function FMINCON); 2) Multi-Agent distributed optimization with ADMM, and 3) Multi-Agent optimization with a GA. Ten runs were made for each operating condition and method to account for the variability of the results which is caused by either the stochastic nature of the GA or the randomly generated initial guess used for the other methods. The average results for each operating condition are shown in Table 2. The results include average power savings (kW), average computational time (s), and the RMS error (kW) in the savings calculations relative to "true" savings from Table 1. The RMSE and computational time results can be more easily visualized in the bubble chart presented in Figure 4.

Table 2. Optimization results comparison

Operating conditions			FMINCON			ADMM			GA		
Test Nr.	Total Load Ton	Wet bulb °F	RMSE Power kW	Average Savings kW	Average time s	RMSE Power kW	Average Savings kW	Average time s	RMSE Power kW	Average Savings kW	Average time s
1	2000	50	15.04	101.55	6.43	0.02	101.5	28.25	0.87	115.7	6.46
2	2000	80	19.83	28.62	6.90	16.08	28.6	86.29	0.02	48.4	6.41
3	4000	50	7.61	326.74	4.44	7.88	326.7	43.55	0.28	334.1	8.30
4	4000	80	17.61	71.89	7.34	25.75	71.9	79.70	0.05	89.4	13.61

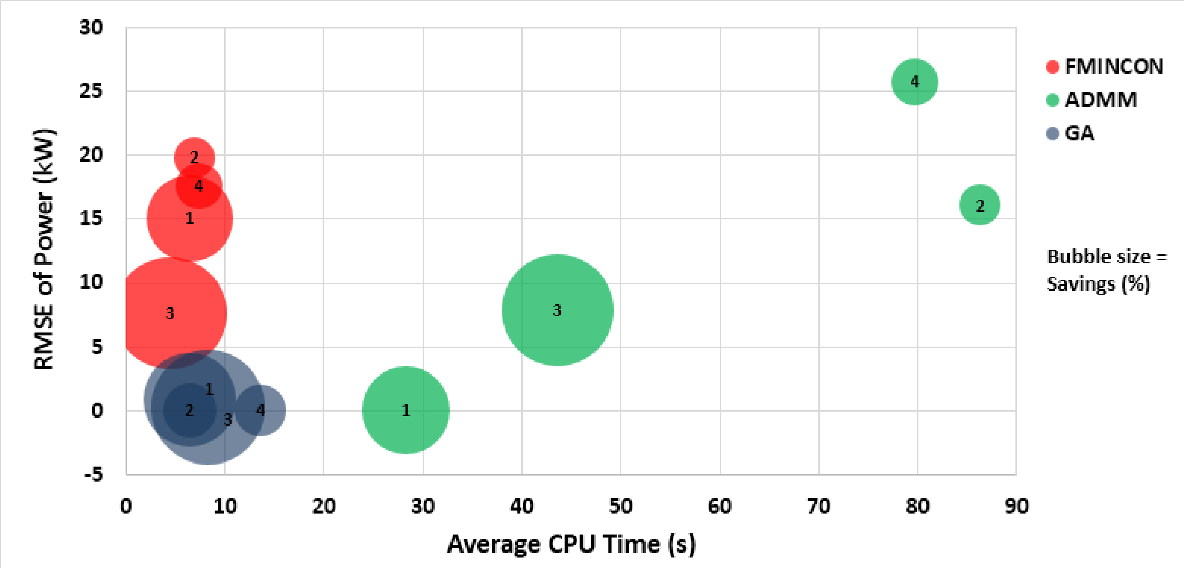

Figure 4. Optimization result comparison

It is important to note that, since the centralized optimization based on mathematical programming techniques (FMINCON) and the ADMM cannot handle discontinuous variables, the optimum number of active chillers for each operating condition had to be specified as an input for these methods. The GA, in contrast, can handle a mixture of integer and continuous variables and the number of active chillers was considered as another optimization variable.

The results show that the GA was able to find near-optimal solutions under all the operating conditions considered. This can be noted in Figure 4, where the RMSE of power (root mean square error in the power consumption compared to the true optimal) was much lower for the GA than the one obtained with the other methods. Further, the processing time for the GA was only 15% of the ADMM (average 8.7s compared to 59.5s). Given the non-convex shape of the cost function, the effectiveness to reach the optimum point and convergence speed of both the conventional centralized optimization and the ADMM method were highly dependent on the operating conditions considered and the goodness of the initial guess.

Comparison with the heuristic strategy, shows that significant power savings can be attained with all the optimization methods. The differences in the predicted average savings are not as significant as might have been expected, except for the second operating condition where the savings obtained with the ADMM and the central optimization with FMINCON are 60% of the savings obtained with the GA.

CONCLUSIONS

In this paper, a multi-agent control approach was applied to a central cooling system without significant storage, such that quasi-steady behavior could be assumed. The resulting non-convex optimization problem was solved for different operating conditions of the plant using a multi-agent simulation framework with two different algorithms: distributed optimization with the alternating direction of multipliers (ADMM) and centralized optimization with a genetic algorithm (GA). Comparison of the results with the "true optimum" point obtained from performance maps of the plant and centralized optimization using mathematical programming techniques showed that the GA was able to find near-optimal solutions for all the cases considered, while the effectiveness of other methods was highly dependent on the operating conditions and the goodness of the initial guess provided for the optimization. Other advantages of the GA are that it can handle a mixture of continuous and discontinuous variables and its convergence speed was much faster than the ADMM.

Even though the GA outperformed the other optimization methods considered in this case-study, the results might be different for a more complex scenario involving a higher number of optimization variables. Further, since the convergence of the ADMM is highly dependent on the goodness of the initial guess it can be expected that in a real setting, the initial guess provided by the optimum point of the previous time step will be closer to the actual optimum point, and the convergence of the ADMM method will be considerably improved. Future work will include an extensive simulation evaluation of the approach using a more complex model of the cooling system and one year of performance data.

NOMENCLATURE

Acronyms

ADMM	Alternating direction method of multipliers
GA	Genetic algorithm
HVAC	Heat, ventilation and air conditioning
RMSE	Root mean square error

Symbols

A	Matrix of coefficients of linear inequality constraints
A_{eq}	Matrix of coefficients of linear equality constraints
Airflow	Cooling tower total air flow
B	Vector of constants terms for linear inequality constraints
B_{eq}	Vector of constants terms for linear equality constraints
C	Feasible region of optimization variables
f	Cost function to be optimized
F_i	Matrix that assigns the local variable array Xi its corresponding global variables
g	Non-linear inequality constraint
h	Non-linear equality constraint
L	Augmented Lagrangian
$\dot{m}$	Mass flow rate
Power	Power consumption
$\dot{Q}$	Heat transfer rate
T	Temperature
T_{mains}	Cooling tower make-up water temperature
X	Local optimization variables
Y	Lagrange multipliers
Z	Global optimization variables

Subscripts/Superscripts

co	Chiller's condenser
ct	Cooling tower cell
cw	Cold well (condenser water reservoir)
db	Dry-bulb
ev	Chiller's evaporator
i	Inlet conditions
i	Subproblem number
k	Iteration number
N	Number of subproblems for multi-agent optimization
o	Outlet conditions
w	Water
wb	Wet-bulb

REFERENCES

Ahn, B. C., & Mitchell, J. W. (2001). Optimal control development for chilled water plants using a quadratic representation. Energy and Buildings, 33, 371-378.

Braun, J. E., Klein, S. A., Mitchell, J. W., & Beckman W. A. (1989a). Methodologies for optimal control of chilled water systems without storage. ASHRAE Transactions, 95(1), 652-662.

Braun, J. E., Klein, S. A., Mitchell, J. W., & Beckman, W. A. (1989b). Applications of Optimal Control to Chilled Water Systems without Storage. ASHRAE Transactions, 95(1), 663-674.

Cai, J., Kim, D., Jaramillo, R. C., Braun, J. E., & Hu, J. (2015). A general multi-agent control approach for building energy system optimization. Manuscript submitted for publication.

Jaramillo, R. C., Braun, J. E., & Horton, W. T. (2014). Application of Near-Optimal Tower Control and Free Cooling on the Condenser Water Side for Optimization of Central Cooling Systems. Third International High Performance Buildings Conference. West Lafayette, IN.

Kelly, G. E., & Bushby, S. T. (2012). Are intelligent agents the key to optimizing building HVAC system performance? HVAC& R Research, 18(4), 750-759.

MATLAB (Version R2013a) [Computer software]. Natick, MA: The MathWorks, Inc.

Summers, T. H., & Lygeros, J. (2012). Distributed model predictive consensus via the alternating direction method of multipliers. In 2012 50th Annual Allerton Conference on Communication, Control, and Computing (Allerton),, 79-84.

Sun, B., Luh, P. B., Jia, Q. S., Jiang, Z., Wang, F., & Song, C. (2010). An integrated control of shading blinds, natural ventilation, and HVAC systems for energy saving and human comfort. In Automation Science and Engineering (CASE), 2010 IEEE Conference on, 7-14.

Torzhkov, A., Sharma, P., Li, C., Toso, R., & Chakraborty, A. (2010). Chiller Plant Optimization: An Integrated Approach for Chiller Sequencing and Control. 49th IEEE Conference on Decision and Control. 15-1.

Treado, S. J. (2010). An Agent-Based Methodology for Optimizing Building HVAC System Performance, ASHRAE Transactions, 116(2), 124-134.

Yao, Y., Lian, Z., Hou, Z., & Zhou, X. (2004). Optimal operation of a large cooling system based on an empirical model. Applied Thermal Engineering, 24(16) pp. 2303–2321.

Zhang, Z., & Turner, W. (2012). Method for Estimating Energy Savings Potential of Chilled-Water Plant Retro-Commissioning. ASHRAE Transactions, 118(2).

Zhu, J., Zhou, D., Li, F., & Fu, T. (2014). Improved Real Coded Genetic Algorithm and Its Simulation. Journal of Software, 9(2), 389-397.

ACKNOWLEDGMENTS

The authors would like to acknowledge the collaboration provided by Purdue Physical Facilities, West Lafayette, IN.