

Software-as-a-service optimal scheduling of New Mexico buildings

C. Birk Jones, Matthew Robinson, Hans Barsun,
Leila Ghanbari & Andrea Mammoli
Mechanical Engineering
University of New Mexico MSC01-1150
Albuquerque NM 87131
USA
cbjones@unm.edu
mr369@unm.edu
hbarsun@unm.edu
lghanbari@unm.edu
mammoli@unm.edu

Salman Mashayekh, Wei Feng
& Chris Marnay (corresponding author)
China Energy Group & Grid Integration Group
Lawrence Berkeley National Laboratory
1 Cyclotron Rd MS 90-1021 & 90-2121
Berkeley CA 94720
USA
SMashayekh@lbl.gov
WeiFeng@lbl.gov
ChrisMarnay@lbl.gov

Keywords

operation, software, smart grid, storage, control, cooling, software-as-a-service, storage state-of-charge

Abstract

The University of New Mexico (UNM) and Lawrence Berkeley National Laboratory (LBNL) are demonstrating building and microgrid remote optimisation using a software-as-a-service (SaaS) configuration for the complex and highly efficient UNM Mechanical Engineering Building (MEBlg), which uses only about 40 % as much energy as typical. The SaaS approach lowers weekly electricity bills by a third under the most favourable conditions, and over an 8week test during summer 2014, an 11 % cost reduction was achieved overall. The building's heating, ventilation, and air conditioning (HVAC) system incorporates cooling assisted by a 232 m² solar thermal array providing heat to a 70 kW_{th} absorption chiller, as well as wintertime heating. A 30 m³ hot thermal storage tank makes heat available at night for both heating and absorption cooling. Additionally, 350 m³ of chilled water (CHW) storage shifts the considerable cooling electrical load of this high desert location off-peak. The big energy and cost savings come from more efficient use of storage compared to the baseline strategy of fully charging the CHW tanks whenever solar production allows. Parasitic loads become significant at these very low energy consumption levels. A MEBlg model has been built on LBNL's Distributed Energy Resources Customer Adoption Model (DER-CAM) platform, and a direct MySQL interface delivers daily week-ahead scheduling based on weather forecasts, loads, tariffs, etc. The approach is complex, involving multiple vulnerable interfaces and execution steps. Failure analysis shows that despite the notable potential cost and energy savings, DER-CAM schedules

are reliably delivered less than half the time, and implementation accuracy for storage charging-discharging and absorption chiller operation is poor, suggesting better methods and device modelling could improve results. The goal of this work is to find low cost methods for achieving bill and energy savings, and not to achieve outstanding technical performance, which might be cost prohibitive given the small absolute energy savings. MEBlg results are reported and set-up described for similar control of a large nearby office building, One Sun Plaza.

Introduction

BACKGROUND

Primarily focused on on-site power generation and microgrid applications, Lawrence Berkeley National Laboratory (LBNL) has been developing building optimising capability for over a decade (Marnay et al. 2008, Marnay and Lai 2012, Marnay et al. 2013). This effort has led to the Distributed Energy Resources Customer Adoption Model (DER-CAM), which provides both optimal equipment choices and operating schedules for commercial building-scale energy systems (Cardoso et al. 2013, Steen et al. 2014). This optimisation engine could be deployed in multiple ways, as a stand-alone on-site scheduler, embedded in a building energy management and control system (EMCS) or individual energy producing, consuming, or storing device controls, or by data exchange between a local control system and a remote DER-CAM server. This latter approach constitutes a software-as-a-service (SaaS) paradigm, and is the structure used in this demonstration. Such an arrangement promises to provide optimisation capability at low cost, and sidesteps the most formidable cost, compatibility, maintenance, and li-

censing hurdles of the alternatives. In general, sophisticated optimisation is most likely to deliver the most noticeable benefit to already efficient buildings with complex energy systems (Cardoso et al. 2013).

The work reported here is primarily an application of the Operations DER-CAM version, which dispatches daily rolling week-ahead optimal schedules to the Mechanical Engineering Building (MEBlg) on the University of New Mexico (UNM) campus in Albuquerque, U.S.A. LBNL-UNM collaboration on this demonstration has continued for three years, and the approach is now being deployed at additional sites; One Sun Plaza (OSP), a large Albuquerque office building, is the second site now under DER-CAM control. These two buildings' efficient but complex systems make them good candidates for SaaS optimisation. The MEBlg was designed and commissioned as a showcase for efficiency and thermal storage, and without DER-CAM optimisation it uses only approximately 40 % as much energy as a typical Albuquerque cohort. Its multiple cooling options together with need for storage charging and discharging pose the type of difficult optimisation problem that cannot be easily solved by intuition or simple rules of thumb. In fact, wherever storage is involved, DER-CAM will likely deliver efficiency improvement because unlike simple operations problems, storage charging and discharging in any timestep affects options in all other timesteps (DeForest et al. 2014). Intuition and simple search algorithms typically cannot perform as well on such tasks as analytic methods, although there may be instances in which the physical characteristics of the system cannot be well represented in an analytic model.

The heart of the technical challenge of a SaaS approach is establishing a robust and secure communication path between DER-CAM and the devices that must follow its optimal week-ahead schedules, which often involves multiple difficult, potentially costly, and unreliable interfaces. The approach is evaluated, and analysis of data delivery success and accuracy of implementation reported. The first key reliability measure is that if schedule delivery fails, the previously received schedule is used, and this proves an effective strategy. Overall, cyber security has been particularly challenging, in part because a high profile December 2013 data breach at Target Corporation took advantage of network access provided for HVAC control. The attack compromised up to 110 million customers' information, and contributed to Target's more than 40 % profit decline in that quarter compared to the same period in 2012. Much of this paper is dedicated to these communications issues, as it is the area in which most progress has been made in recent work.

The work reported here has been primarily financed by the U.S. Department of Energy, Building Technology Office via the U.S.-China Clean Energy Research Center on Building Energy Efficiency (CERC-BEE) program. The approach will be demonstrated in China, and also by the local electricity utility, PNM. The overarching program goal is achieving ultra high efficient buildings and microgrids with local generation and control. Nonetheless, the operations approach is economic, with optimisation typically driven, as in this work, by cost minimisation, which is likely to stimulate widespread deployment, although energy or carbon footprint minimisation is readily implementable alternative objectives. Consequently, the second experiment reported herein focuses on cost savings achieved

during an 8-week summer 2014 week-on-week-off test period. Further, note that when buildings and mechanical systems are already highly efficient, yielding significant cost savings based on energy efficiency becomes increasingly difficult, while reducing costs by taking advantage of tariff structures may be far more effective. Not addressed in this work, but looking beyond tariffs, tight systems control can enable trading in energy, demand response, and ancillary services markets to generate additional smart grid revenue streams.

UNM MECHANICAL ENGINEERING BUILDING DESCRIPTION

The MEBlg, shown at left in Figure 1, is a 7,000 m² building energy systems *living laboratory*. The building was commissioned in 1980 with complex thermal (both hot and cold) storage capability, which received a thorough modernisation between 2006 and 2010 (Ortiz et al. 2010, Mammoli et al. 2010). The HVAC system also incorporates cooling assisted by a 232 m² solar thermal array with an extended peak enabled by booster mirrors, a 70 kW_{th} absorption chiller, and virtually all solar wintertime heating (Armenta et al. 2011).

A campus district energy system (DES) supplies electricity, CHW, and steam to all buildings on campus. During the cooling season, mid-March to mid-October, the building can be cooled by its large thermal cold storage, a solar-powered absorption chiller, the campus DES CHW, or by a dynamic combination of the above. This level of complexity presents an ideal application for schedule optimisation as executed by DER-CAM.

ONE-SUN PLAZA BUILDING DESCRIPTION

OSP, shown at right in Figure 1, is a commercial office complex commissioned in 2001. The complex is actually two buildings, San Francisco and Pan American, both served by a 1,500 m³ stratified CHW storage system and two 1.1 MW_{th} electric chillers. Despite its apparent simplicity, OSP offers an interesting optimisation opportunity combined with commercial implementation barriers, notably security and cost.

Methodology

As noted above, this project had two primary goals, namely to evaluate the effectiveness of the SaaS approach in this application, and to demonstrate savings in summertime operations costs. Scheduling savings costs can result from improved energy efficiency, and also from better exploitation of tariff structures. Also possible, although not covered in this work, are energy, demand response, ancillary services, or other market participation, additional revenue opportunities requiring effective scheduling. In general, savings from both efficiency and tariff response are observed, but this is not always the case. While the ratio of savings from the two mechanisms can differ substantially in magnitude and sign as a function of local conditions, particularly tariff structures, data from this experiment will be used to obtain a better understanding of the cost drivers and if possible to make some generalisations useful for the design of future microgrids. Locally in Albuquerque, the tariff has both time-of-use energy charges and a power (demand) charge (PNM 2011, and shown in Table 2). The on-peak period is 8:00–20:00 weekdays. Such tariff structures tend to make the tariff the key driver among inputs. The experience of a practical implementation of the SaaS service on a research

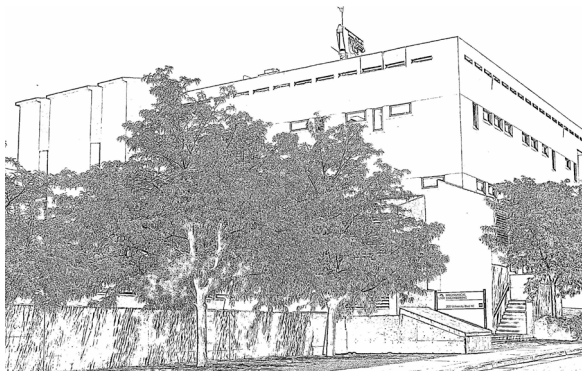
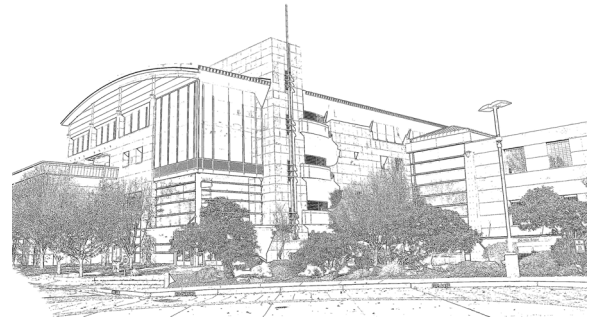


Figure 1. UNM MEBldg, left, and OSP, right.



testbed and on a commercial facility has also resulted in a better understanding of both the benefits of the approach and the potential difficulties.

A data acquisition system was implemented, using infrastructure independent of any proprietary EMCS system. For the MEBldg, essentially all control and monitoring data are recorded and archived, typically at 5 min resolution, or on change of value, as appropriate. For OSP, only data relevant to the present task are recorded and archived in a MySQL database, for example, flow rates, tank temperatures, and temperatures at strategic locations. In addition, data on the performance of various technologies, e.g. storage tank losses, pump power consumption, etc., were obtained either by direct measurement or from the technical literature (Mammoli et al. 2013).

To truly understand the benefits of using DER-CAM even on facilities that are relatively small and with relatively simple Distributed Energy Resources, it is important to understand the sources of cost savings and what erodes them. For example, thermal diffusion from the warm layer into the cold layer of a stratified cold storage tank results in reduced cooling capacity in the air handling units (AHUs), consequently increasing fan power consumption or even activating expensive on-peak electric cooling. In addition, residual cold water left overnight has a higher potential to warm up due to conduction through imperfectly insulated concrete walls. In this work, the interaction of thermal and electric components was considered carefully, e.g. the value of operating an absorption chiller at night, avoiding costly on-peak operation of parasitic loads. Lower nighttime ambient temperatures also improve chiller efficiency, although this is not considered in the current optimisation.

Conceptually, in highly efficient building systems and microgrids, a trade-off must be made between complexity and cost savings from efficiency. Increasing complexity has diminishing cost returns, and also increased probability of failure, as well as a higher initial implementation cost, as illustrated in Figure 2. The analytic problem indirectly being addressed in this program is what optimal level of complexity provides the best compromise of cost savings versus acceptable deployment cost and failure risk. In Figure 2, as complexity increases along the x-axis, implementation costs (solid curve) increase, while savings (dashed curve) are subject to diminishing returns, and become increasingly difficult to generate. While these effects are widely recognised, the importance of reliability (dotted curve) much less so. As complexity increases,

the risk of failure increases and reliability erodes. The overall cost (shaded area) is therefore uncertain, and this must also be considered in selection of the optimal complexity level, which consequently lies to the left of the strict minimum cost position.

More specifically in these buildings, demand charges mean the risk of failure of one element of the SaaS optimisation data flow can eliminate cost savings accrued over a billing cycle. This could easily happen, for example, if storage is not sufficient to provide the afternoon cooling requirements and backup systems come online during a peak tariff period. The ratcheted nature of demand charges make cost minimising scheduling a very complex optimisation problem that remains to be solved. In this work, the monthly peak is always assumed to fall within the rolling 7-day-ahead scheduling horizon.

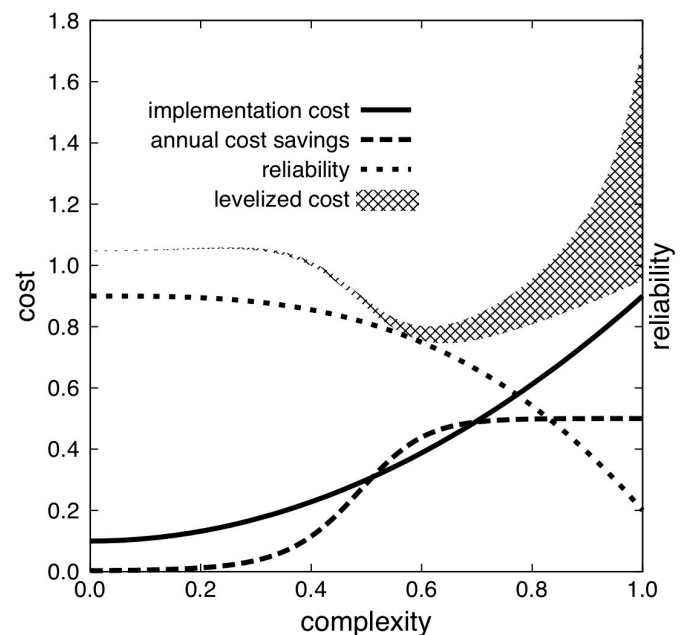


Figure 2. Interaction between optimisation cost savings, implementation costs, reliability, and levelised energy costs as a function of system complexity. Note that this is a qualitative representation, rather than a plot of actual data.

Experiment equipment

UNM MEBlg MECHANICAL ENGINEERING BUILDING

The MEBlg was designed to be a testbed for building-scale DERs, combining solar thermal components and thermal storage. A four-stage centrifugal pump (P1 in Figure 3) driven by a variable speed motor circulates a glycol-water mixture through the solar array. A second pump (P2), also driven by a variable speed motor, draws water from the bottom of the hot storage tank, routes it through the heat exchanger (X1), and returns it to the top of the hot storage tank. The solar field produces $170 \text{ kW}_{\text{th}}$ at peak, with the extended booster mirror peak period, well in excess of the $100 \text{ kW}_{\text{th}}$ needed to drive the absorption chiller's regeneration. This heat accumulates in the 30 m^3 of hot water storage, while excess CHW is stored in the 350 m^3 of cold storage. At night, the cold storage is charged by the campus DES via heat exchanger X2. This storage consists of highly-insulated, underground, unpressurised concrete tanks.

In summer and shoulder season operation, the rooftop solar thermal array charges the hot water storage tank, via the brazed plate heat exchanger X1. The hot layer is maintained at approximately 82°C , while the cooler layer is at 71°C . Hot water from storage can be dispatched by pump P3 to the absorption chiller, operated so the temperature difference between hot water inlet and outlet is maximised, corresponding to lowest heat medium flow rate allowable. In turn, the chiller draws water from the air handler return and from the cold storage, at a mixture temperature of approximately 15°C and cools it to about 7.5°C . Cold water from the chiller is routed to the AHUs. Any portion of the CHW production not needed by the air handler is returned to the cold storage. Low-grade waste heat from the chiller is removed by the cooling tower. The primary source of energy for the absorption chiller is zero operating cost solar heat; nonetheless, several pumps are required to operate: the 0.3 kW

chiller, 0.8 kW internal solution, a 0.5 kW CHW circulation P6, and 3.4 kW cooling tower P4 pumps.

In this experiment, the control variables chosen for optimisation were the chiller run time and the cold water tank charge and discharge profiles. In practice, actively controlling charging and discharging is very challenging because of the multiple automatic systems that control various processes. Intuitively, the absorption chiller would run off-peak, to minimise electricity costs; however, the capacity of the hot storage tank is insufficient to hold the entire amount of heat produced by the solar array on a clear day, i.e. the chiller must run, at least partially, during the on-peak period if the CHW will be needed. A challenge and opportunity of the day-ahead optimisation is to minimise the amount of on-peak operation as a function of weather and cooling load. A second opportunity for optimisation arises from the amount of charge provided to the cold storage tanks to meet the following day's building cooling load. The MEBlg tanks were sized to completely satisfy demand on the hottest day of the year. As a consequence, a full charge is almost never needed. Overcharging during a period results in excessive pump energy use and heat losses during the following night. The optimisation opportunity tends to be estimating how much to charge the tanks so that minimal or no charge remains at the end of each day. Note both that optimisation rests on small loads and losses often ignored in scheduling.

ONE-SUN PLAZA BUILDING

The OSP cooling load is also served in several ways, as shown in Figure 4. If the tanks are at least partially charged, CHW is drawn from the bottom of the tank and routed to a set of low-pressure cooling coils in associated AHUs. In addition, water flows to the heat exchanger, which serves a high-pressure (400 kPa) loop feeding a set of zone-specific fan coil units. It is also possible to operate the electric chillers in parallel with the tanks. In this case, CHW produced by the electric chillers can supplement

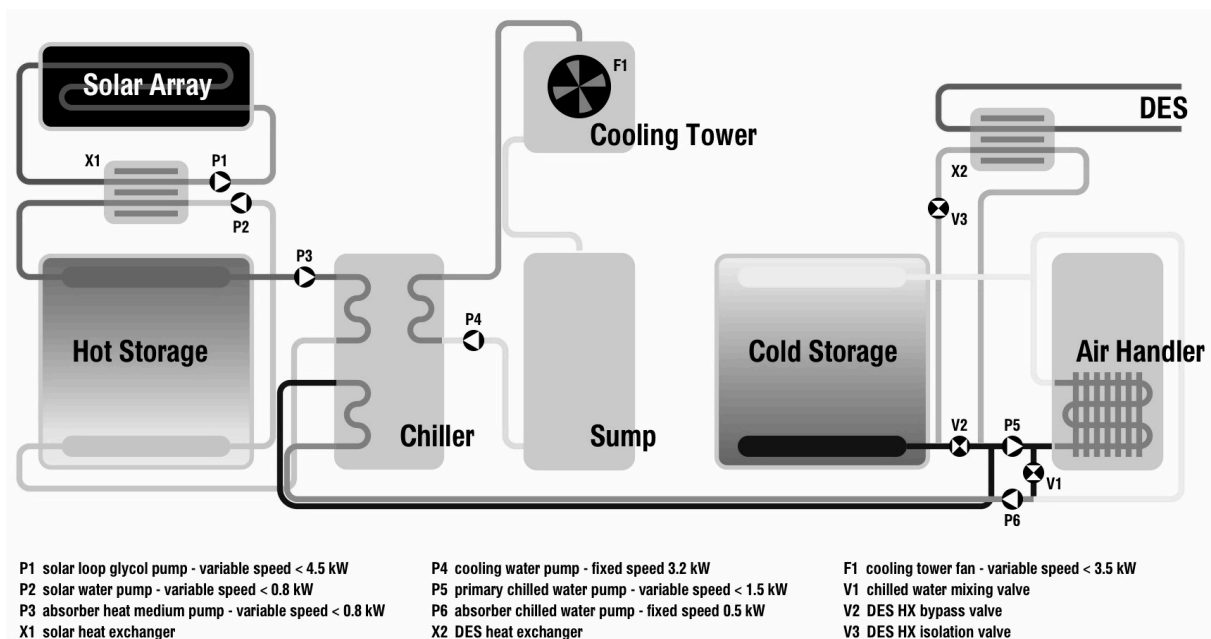


Figure 3. MEBlg thermal system that includes a solar array, hot and cold storage tanks, absorption chiller, and air handling units.

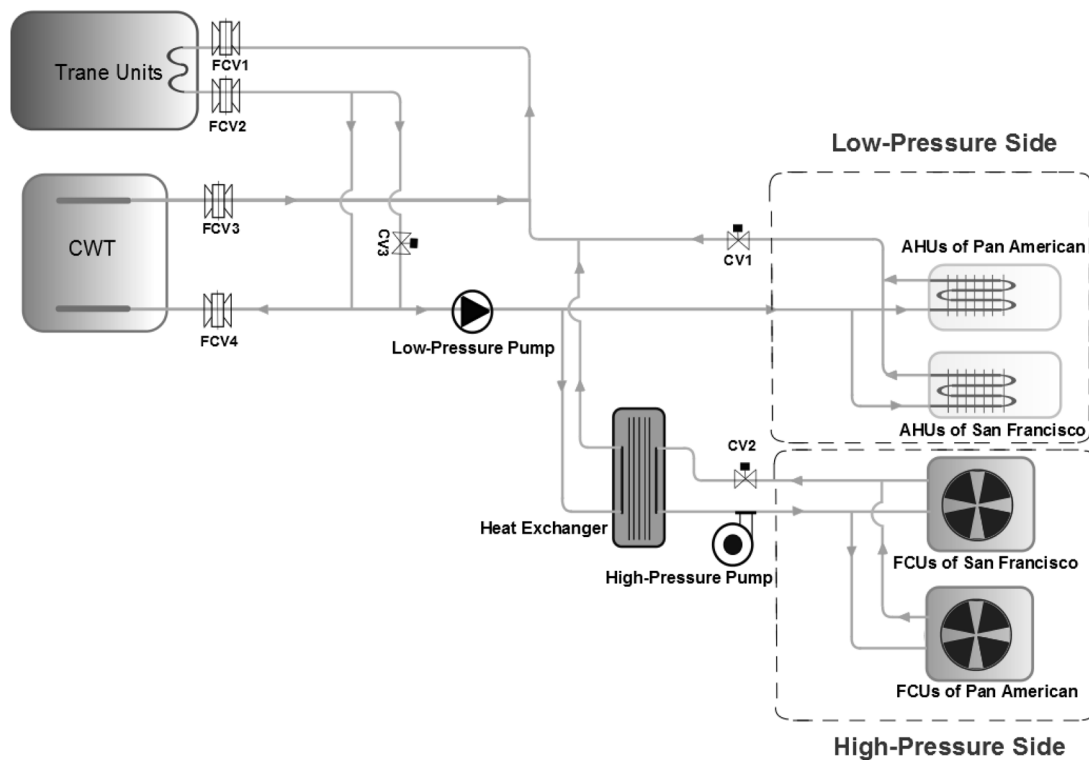


Figure 4. OSP thermal system that includes chillers, heat exchanger, CHW tank, and air handling units.

water supply from the tanks, or even recharge the tanks while cooling the building, if the building cooling load is low.

As with MEBldg, opportunities for cost saving come from both efficiency and exploitation of the tariff structure. From the point of view of efficiency, the strategy is to charge the cold storage tanks only as needed, to avoid the situation where residual charge remains at night. There is one significant difference between OSP and MEBldg. While the latter is designed so that cold storage can always meet the total daily load, storage in the former facility can only partially meet the load on hot summer days. Consequently, it is necessary to operate the chillers for some time during the on-peak period. This provides an opportunity to optimise when and how the tanks are depleted, in such a way that peak demand is limited, an operational task that would be challenging, even for an experienced facility operator. In general, DER-CAM is calculating a timestep by timestep state of charge (SOC) for storage. Unfortunately, setting these SOC values as dynamic targets proves difficult in practice because so many other systems limit the capacity of storage charging and discharging through time. Consequently, in this project only total nighttime charging is optimised at both buildings.

System infrastructure

DATA FLOW OVERVIEW

Centralised optimisation for on-site generation and microgrids provided by DER-CAM, seen top right in Figure 5, must be distributed to local facilities through a network of information technology (IT) connections. In the case of the two demonstration buildings, UNM effectively acted as an energy service provider, or a Distributed Energy Resource Customer Integra-

tion Provider (DER-CIP), shown bottom right of Figure 5. UNM developed, managed, and implemented load forecasting, data transfer, and integration with system controls. The initial implementation of the optimised schedule service from DER-CAM was achieved through a File Transfer Protocol connection and transferred data using Excel files (Mammoli 2013). Recent updates to the system use a MySQL database with Python and PHP scripts.

Figure 5 provides a detailed overview of the IT infrastructure and automated data transfer process performed by Python and PHP: Hypertext Preprocessor (PHP) scripting. Each day the following 7 tasks are executed. 1. Software in the DER-CIP framework extracts the 7-day-ahead weather forecast from the U.S. National Oceanographic and Atmospheric Administration web-site. The load forecast tools use the weather data as inputs and calculate the 7-day-ahead load for the respective buildings. 2. Forecast results were archived in the MySQL database. 3. DER-CAM extracts this load forecast using an Excel interface. 4. Using these inputs, DER-CAM executes the optimisation based on its mixed integer linear formulation written in the Generalized Algebraic Modeling System. 5. The optimised schedule is routed back to the Excel interface. 6. Excel reinserts data into the MySQL database. 7. The schedule is routed to the respective buildings. Establishing, trouble shooting, and maintaining this final connection to the two buildings was a critical task performed differently in order to adapt to the missions, resources, and constraints of each site.

The connection to the two buildings considered reliability and cyber-security. The connection between the MEBldg and the central database applies PHP code utilising the Open Database Connectivity connection. This code queries the database for the optimised schedule and then inserts values into analog variable

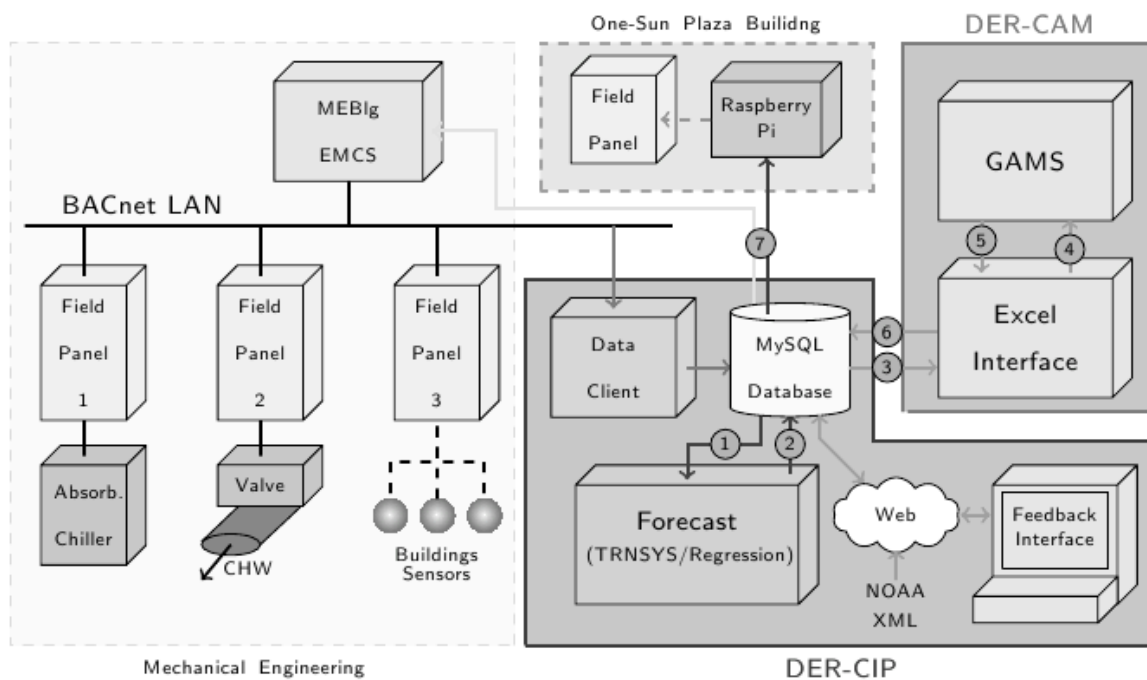


Figure 5. Data transfer process and infrastructure that connects the optimised schedule calculation from DER-CAM to the EMCS at MEBlg and a Raspberry Pi at OSP.

points in the proprietary control system managing the absorption chiller and DES CHW inflow valve. In contrast, the connection with OSP is achieved through SQL queries to a Raspberry Pi web-server, which is situated on a network available to a single remote IP address not connected to any network inside the building. This arrangement was necessitated by the cyber-security concerns of the building owner. Schedule data are routed to the Raspberry Pi and stored locally in a MySQL database. Then, a Python script interprets the schedule and provides an electrical signal to an analog input to a field panel. The amount of voltage provided by the electrical signal defines the amount of CHW charging required to meet the next day's thermal load. Naturally, all of the interfaces described are costly to establish and maintain, as well as vulnerable to failure; hence, the great interest in this work in evaluating them and developing more robust structures. Also note that analog interfaces for low bandwidth data exchange have significant cyber security advantages.

FORECASTING

This paper reports on two options for developing the 7-day-ahead operations load forecast needed by DER-CAM. Load forecast at the MEBlg were performed using a Transient System Simulation Tool (TRNSYS) model that was run automatically each day. The load forecast at OSP was based on a regression model. This regression model was fitted to TRNSYS simulation results. The calibration process for the MEBlg TRNSYS model used standard techniques, and the OSP model implemented a genetic algorithm (GA).

The TRNSYS model that emulated the MEBlg was calibrated by minimizing the difference between the prediction and the measured load. This comparison complies with ASHRAE Guidelines 2002 whereby the Normalised Mean Bias and Coefficient of Variation of the Root Mean Squared Errors must be

below 10 % and 30 %, respectively. After about ten iterations of internal loads (occupancy, computer, etc.) being adjusted, compliance of both error criteria was achieved.

A TRNSYS model of the MEBlg existed from prior work, but the construction of the OSP TRNSYS model implemented a process that discovered and optimised parameters to quickly reduce the error between actual and model results (Jones 2010). The process begins with a sensitivity analysis to identify the most influential parameters causing significant changes to the cooling load. Then, a genetic algorithm from the MATLAB toolbox was implemented to find the best values for defined parameters. This was accomplished by running the TRNSYS model several times, and at each iteration the input parameters are updated based on values provided by the GA. The iterations continue until convergence defined by the actual and modeled error or when genetic algorithm outputs stabilize.

BACNET SENSOR NETWORK

Feedback from the actual MEBlg equipment was acquired through a data client that accessed sensor data in the EMCS, which was possible because it uses the standard BACnet protocol. It has been approved by ASHRAE, ANSI, and ISO as the communication standard for all commercial buildings. BACnet provides reliable communication networks for controls and sensors of HVAC, lighting, access control, and fire detection systems. All new commercial buildings in the U.S. must comply with this standard, and many existing buildings already do.

According to the 2012 Commercial Building Energy Consumption Survey there are about 5.6 million commercial buildings in the United States, over half of which have floor space over 4,650 m². Furthermore, it is estimated that 70 % of these buildings have an EMCS (Kiliccote et al. 2006). This suggests that there is a significant potential for penetration of the demonstrat-

ed approach into the existing building stock; however, there are significant barriers to overcome for complete and reliable integration of a SaaS remote scheduling service into existing control systems. Certainly, communication via standardised protocols such as BACnet or OpenDR can facilitate faster deployment.

INTERFACE

The DER-CAM based SaaS has the potential to provide a valuable, cost saving service to commercial building managers; however, its implementation faces significant obstacles. In this work an additional layer was developed to provide an interface between DER-CAM and the EMCS. Also, feedback from the building and the optimisation is provided to the end-user via a web-based interface. Feedback is essential to provide end-users (facility managers) with confidence the service is operating correctly and beneficially. This platform could promote increased adoption of such services, which could be provided on a performance contract basis. Potential benefits of the SaaS approach result in part from associating the various responsibilities to specialised groups with appropriate skill sets. The activities can broadly be categorised into three major groups: 1) developing and maintaining the optimisation service; 2) building and maintaining interfaces between the facility and the optimisation service; 3) providing historical and operational data and experience. The structure of the SaaS application is concentric: a centralised service of highly trained experts focuses on the core optimisation service, in this case DER-CAM. The second circle is a set of energy service providers, distributed regionally, who are responsible for connecting the optimisation service to the end user. Personnel involved in this activity would likely be composed of engineers with experience in relevant modeling methods and IT, as well as with day-to-day building operations. The third ring consists of the end users, and specifically facility managers, with long-term close experience with their specific facility. This specialisation structure ensures that all people involved spend the majority of the time on tasks related to their skill set, maximising the efficiency and quality of the end result. In contrast, managing complex optimisation algorithms is not easily achieved outside a laboratory setting, and would pose a challenge to facility managers in the outer ring.

The expected implementation cost saving also emanates from this concentric structure. The optimisation hub must ensure continuous operation of the service, provide a well-understood and well-documented interface for the optimisation engine, and offer some customisation for unusual situations. The interface service provider must obtain all the data required to produce weekly forecasts of load, either using historical data (seldom), or by building a sufficiently accurate model of the building (often); further, the provider must process information about the DER technology available at the facility. Finally, the end user's task is to archive and share historical data, equipment specification, architectural/engineering drawings, etc. as needed for modeling, and then facilitate the implementation of optimised scheduling into the building EMCS or other controls.

Experiment

The optimisation service presented in this paper attempts to capitalise on opportunities for minimising on-peak equipment use and reduce over charging of the CHW storage.

These achievements were dependent on the reliability and performance of the service; therefore, reliability and performance were monitored and evaluated over the 8-week period. During this time, which started on 5 Aug 2014 and ended on 6 Oct 2014, the MEBldg absorption chiller and CHW tank systems were controlled by the optimised schedule for four weeks and by normal operations the other interspersed four weeks. This experimental design was implemented in an on/off manner, where the control's settings were changed each Monday between optimisation or baseline mode. While each of the modes was in operation, actual, forecast, and optimisation results were collected and evaluated.

Evaluation of results was broken out into two experiments, with the intent of clearly defining and quantifying system successes and deficiencies. The first experiment focused on the reliability of the SaaS approach. The reliability was dependent on the IT infrastructure and sequencing of operations. Consequently, a review of missed data communications was compiled to define probability of optimal schedule transfer success. The second experiment concentrated on the cost performance of the optimisation over the baseline. This review considered performance metrics such as cost versus load, and optimisation versus actual performance.

Results

EXPERIMENT 1: SAAS RELIABILITY

A key component of the SaaS approach (or any remote service) is the probability that the information provided (in this case, a schedule) makes it through the communication chain's links, from the optimization solver to the building automation system. A previous study of SaaS control of the MEBldg, conducted the summer of 2012, implemented DER-CAM in an error prone fashion. The results were not widely reported, but it was concluded that schedule transfer must be redesigned to improve reliability. Additionally, a metric of reliability was sought so progress could be tracked as process changes were implemented. Only the MEBldg has thus far gained sufficient experience to make such estimates.

Figure 5 illustrates the data transfer path via the connections 1–7. Each attempted data transfer across a connection can be considered a random experiment described as a discrete random variable with a binomial distribution of probability, p , of success, S , i.e. accurate data delivery. For each attempt over the 8-week experiment, 1 or 0 was written to a local MySQL database to represent success or failure, respectively. If a failure occurred, a note was taken that documented location and cause. This showed where work should be concentrated to improve transmission, and these results allow calculation of overall likelihood of successfully receiving a *daily* schedule. The chance of receiving a schedule is equivalent to the simple product of each individual probability of success at each connection, that is,

$$P(S1S2S3S4S5S6S7) \\ = P(S1)P(S2)P(S3)P(S4)P(S5)P(S6)P(S7) \quad (1)$$

where the subscript denotes the 7 connections described above and shown in Figure 5 (Montgomery and Runger 2014). Fig-

Baseline —
Optimized — Schedule Error x

S	M	T	W	T	F	S
Aug						
27	28	29	30	31	1	2
3	4	5	6	7 x	8	9
10 x	11 x	12	13	14	15	16
17 x	18	19	20	21	22 x	23
24 x	25	26	27	28 x	29	30 x
31 x	1	2	3	4	5	6

Figure 6. On/Off schedule used for UNM ME DER-CAM experiment.

Table 1. Probabilities of data connections in Figure 5. N=62.

connection	successes	p(S)
1	61	0.984
2	61	0.984
3	49	0.790
4	62	1.000
5	62	1.000
6	40	0.645
7	62	1.000

	Hour	1	2	3	4
38					
39	Heat from DG	0	0	0	0
40	Heat from NG	0	0	0	0
41	Heat from Solar Thermal	0	0	0	0
42	Heat Load	0	0	0	0
43	Heat Storage SOC	#DIV/0!	#DIV/0!	#DIV/0!	#DIV/0!
44	Heat to Storage	0	0	0	0
45	Heat from Storage	0	0	0	0
46	Heat lost to Decay in Storage Tank				

Figure 7. Data type error that occurs between DER-CAM interface and DER-CIP database, at right above. The #DIV/0! is a string value that the database does not understand and the transfer fails.

ure 6 shows the on/off schedule over the 5 Aug to 6 Oct period, with baseline weeks shown in light type and DER-CAM optimised weeks in bold type. The x indicates an error along the data path illustrated in Figure 5.

During the eight weeks, the success and failure of each connection was monitored to estimate the likelihood of successfully receiving a schedule. Success rates of these 7 links are depicted in Table 1. Assuming failures are independent, the probability of a complete successful schedule transfer on any given day can be estimated using eqn. 1 above and the observed probabilities presented in Table 1, and it is 0.494, less than an even chance. Given DER-CAM provides a daily week-ahead schedule, one might ask what is the likelihood that the equipment will not receive an updated schedule over a given 7-day period? If Y is a discrete random variable that describes the event of updating a schedule *at least once* in seven days, the chance of getting an updated schedule is the complement of Y , seven straight days of failure, and is given by the following.

S	M	T	W	T	F	S
Sep						
31	1	2 x	3 x	4 x	5 x	6 x
7 x	8 x	9	10 x	11 x	12 x	13 x
14 x	15 x	16 x	17 x	18 x	19 x	20 x
21 x	22 x	23	24	25 x	26 x	27 x
28 x	29	30 x	1 x	2 x	3 x	4 x
5 x						

$$P(Y \geq 1 \text{ Success}) = 1 - P(Y = \text{All Failures})$$

$$= 1 - 0.506^7 = 0.992 \quad (2)$$

The above result predicts that over the long-run, about 3 times a year no DER-CAM schedule will arrive for seven days. To mitigate this high failure rate, connections 3 and 6 will be examined. Connection 3 failure was produced by problems with the LBNL server. The switch to a more reliable server after the second crash increased reliability substantially. Connection 6 failure, DER-CAM not transferring data to the DER-CIP database, was caused by a data type mismatch. In certain instances, the optimisation produces a string value in the DER-CAM interface where the DER-CIP database required float data (Figure 7). This problem could be alleviated by adding robustness to the DER-CAM interface, when certain values are returned from Generalized Algebraic Modeling System and forwarded.

It has been shown that the probability of not getting a schedule for the 7-day optimisation horizon is relatively low compared to the chance of getting a new schedule every day. In other words, the week-ahead rolling forecast window significantly improves performance. The worst case scenario, i.e. not updating the schedule for seven days, occurred during the experiment during 9–15 Sep. In this blackout, the DER-CIP was continually updating its forecasts and saving them to the local database, but the last schedule the building received arrived on 9 Sep. The load forecasts should have been sent to DER-CAM every day when the server was down, but could not be delivered. The comparison over that week provides insight into the consequences of the schedule not being updated for over a 7-day horizon.

Figure 8 (left) shows that the difference between the forecast utilised for the 9 Sep week and an updated daily forecast is small. The discrepancy between the actual thermal load and the predicted load, at right, regardless of current or last update, is larger than the difference between the current and optimisation forecasts. In other words, given the uncertainty in load forecasts, the simple expediency of the rolling week-ahead forecast performs reasonably well. As described, the major communication failures resulted from relatively trivial fixable problems. The more complex question shown in Figure 2 concerns the trade-off between expensive fixes and improved communications, which may not be worthwhile in terms of bill savings.

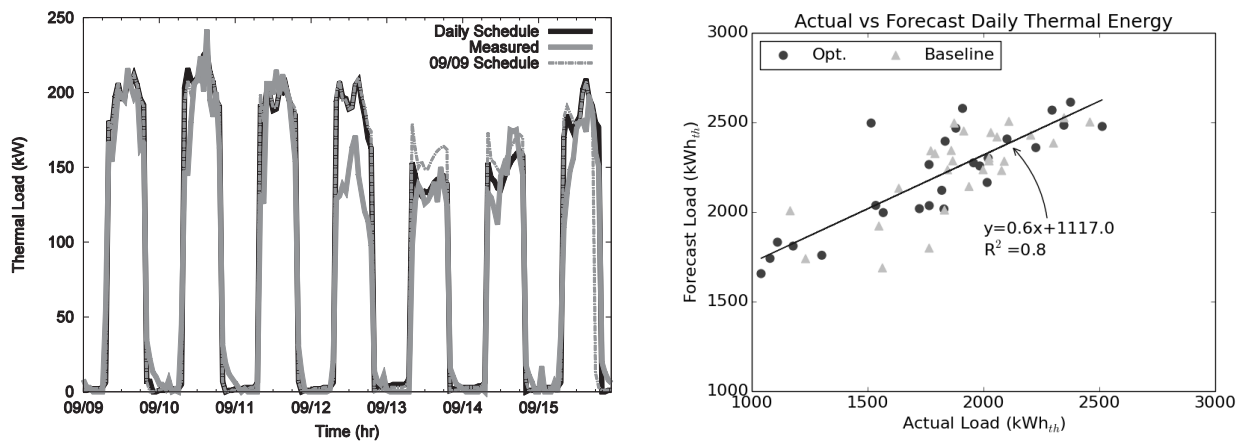


Figure 8. Left: A comparison between the schedule received on 9 Sep. 2014 and the daily updated schedule. Deviations between the current forecast and the last updated schedule are small. Right: MEBlg scatter plot of actual vs. forecast thermal load for the entire experiment. According to the linear fit, the forecast over predicted at low loads, but performed well at high loads.

EXPERIMENT 2: COST SAVINGS

Forecast performance

The MEBlg daily load forecast simulation, described in the Forecasting section above, was run each day and provided an estimate of the cooling requirement for the next seven days. The results, provided in five-minute intervals, were compared with actual measured values to evaluate the tool's accuracy. A graphical representation of this comparison is described in Figure 8 left, above. In this example, the forecast results were consistent with the measured values, and through visual inspection of the discrepancy, it was evident a sufficient fit had been accomplished. Additionally, the calculated Normalised Mean Bias and Coefficient of Variation of the Root Mean Squared Errors, 8 % and 19 % respectively, are well within the ASHRAE Guidelines. Nonetheless, the model did not provide a good representation of measured values collected on Fridays when school was in session. This unfortunate miss is evident on Fri 12 Sep in Figure 8.

The sub-hour comparisons of model versus actual were reviewed to assess model physical property performance. The overall goodness of fit for the model considered the thermal energy results for each day. On a macro-scale, the total TRN-SYS forecast overall total thermal energy requirement for the 8-week period was about 122 MWh_{th} while the actual building use was 103 MWh_{th}, i.e. the model over predicted by about 18 % for the length of this experiment.

The scatter plot, shown in Figure 8 right describes the daily forecast thermal energy results versus load. The bold circles represent the forecast, and the light triangles signify values collected during the baseline periods. The green line shows the least-squares regression considering all data, with an intercept of over 1,100 and a slope of 0.6. This intercept, far from the origin, indicates forecast load over predicted actual thermal energy needed, when the cooling requirement was low, probably due to inaccurate occupancy inputs; however, as the thermal cooling requirement increased, the model became more accurate.

Optimisation performance

Each day DER-CAM produces an optimised 7-day-ahead schedule. Its algorithms use the load and weather forecast to determine the best way to operate the absorption chiller and charge-discharge the CHW tanks. Results from the algorithm were then routed to the central database and used as input control points in the EMCS. As analysed above, this process did not succeed every day, but contingencies were implemented so the next day schedule from the last successfully received 7-day schedule was implemented. As a result, the system was able to provide the appropriate amount of charging at night, avoiding DES CHW use during the on-peak utility pricing period about 93 % of the time; however, this percentage is calculated from a small sample size of 27 days. Additionally, the success rate for baseline operations was also 93 %. This suggests that other issues, such as CHW distribution control malfunctions at the campus CHW system, could have impeded actual off-peak charging to not reach the desired level. Also, note that on-peak usage in itself does not necessarily mean the schedule is suboptimal.

Figure 9 describes an example, from 19–21 Aug, where multiple possibilities could have caused daytime DES CHW use, which occurred on both 21 and 22 Aug. On 19 Aug, optimised CHW production from off-peak charging and the absorption chiller was 2.4 MWh_{th}, yet actual was only about 1.4 MWh_{th}, due to a campus CHW distribution controls malfunction causing off-peak charging to cease around 2:00 h on 20 Aug. But the optimisation results were problematic as well. The optimisation specified a CHW tank SOC of 0.74 and an absorption chiller output of 0.5 MWh_{th} for a load forecast of 2.4 MWh_{th}. It was discovered, on 17 Aug, that for a very similar load forecast and the same absorption chiller output, the optimisation calculated a more realistic SOC of 0.84 rather than 0.74. It then became evident that the 0.74 target SOC occurred about 38 % of the time at varying load forecasts through the experiment.

During the total four-week optimisation period, total CHW tank charging energy from the campus distribution system and the local absorption chiller were 54.8 MWh_{th} and 57.6 MWh_{th}

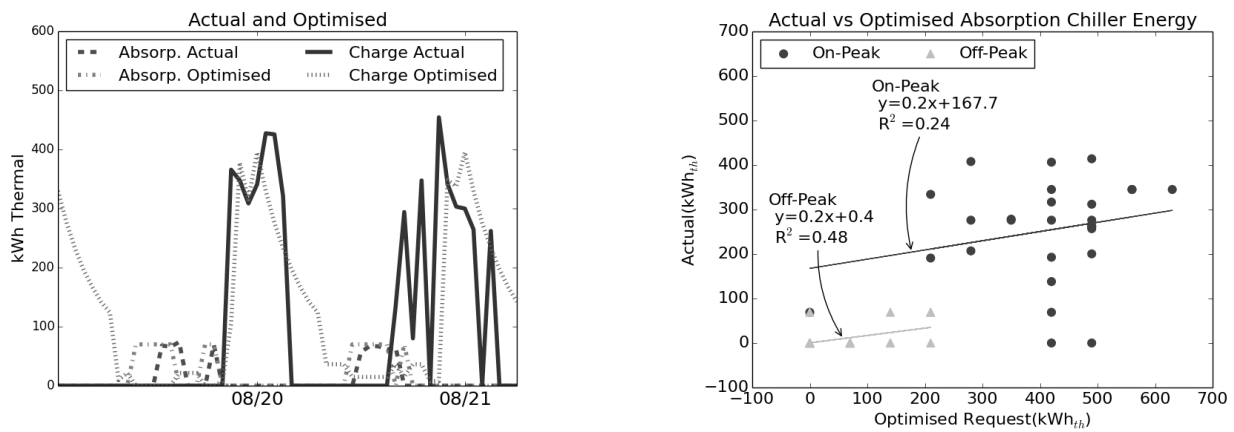


Figure 9. Left: Actual vs. optimised absorption chiller operation and tank charging over an August two-day period. The tank was insufficiently charged requiring daytime chiller operations. Right: Scatter plot of actual and optimised chiller operation shows significant discrepancies between desired and actual performance.

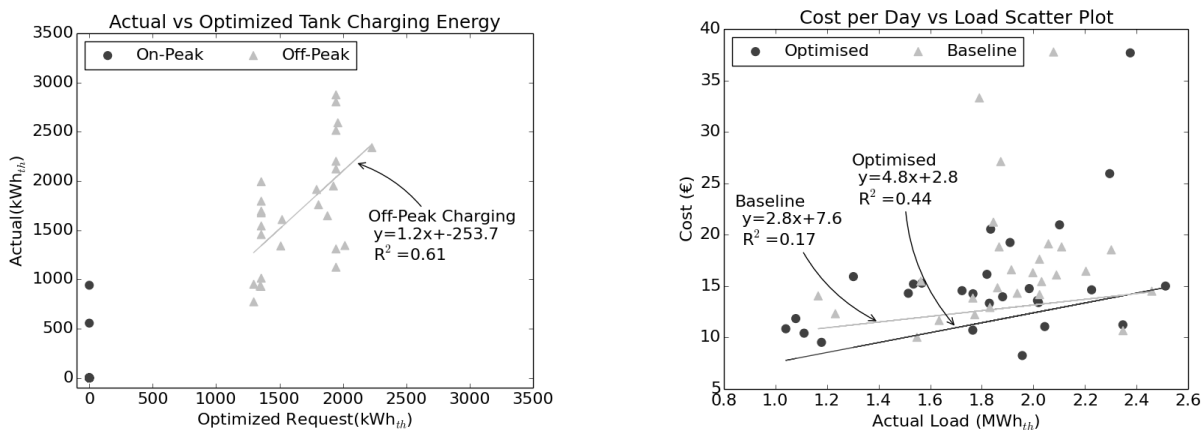


Figure 10. Left: Scatter plot of actual and DER-CAM optimised requests for CHW tank charging. The regression analysis indicates a low correlation possibly due to the optimisation not considering the tank SOC at the end of each day. Right: Scatter plot and regression of the cost per day in relation to the actual thermal load of the building. It is difficult to draw conclusions on the potential cost savings because the regression analysis produced low R^2 values. However, the analysis hints that there is a cost savings potential at low loads that diminished as the load increased.

for the actual and desired, respectively, an overall error of only around 5 %; however, the daily optimised and actual absorption chiller energy production did not match well. Average production during the on-peak was 251 kWh_{th} and 407 kWh_{th} for actual and optimised operations respectively. Additionally, Figure 9 right shows optimised vs. actual off-peak absorption chiller results. The regression lines shown for on and off-peak production indicate low correlation between, with R^2 values of only 0.24 and 0.48, respectively, showing the actual absorption chiller could not produce as much thermal energy as the optimisation demanded. This discrepancy could be attributed to inaccurate prediction of solar irradiance during clouding days, or failure to recognize the chiller's technical limitations, e.g. warm up time.

The average CHW energy to charge the tank during off-peak operations were similar for actual (1.71 MWh_{th}) and desired (1.67 MWh_{th}) cases, although with standard deviations varied at 577 and 298 for the actual and desired, respectively. Figure 10

provides a scatter plot of the actual and optimised results for the CHW tank charging. It shows that the correlation between actual and desired for off-peak results is not significant with an R^2 value of 0.61. The discrepancy could be attributed to the fact the optimisation was unaware of the CHW tank SOC at the end of the day. Therefore, actual results varied at different optimised SOC requests provided by DER-CAM.

Cost comparison

The overall electrical cost incurred during this experiment to support the absorption chiller and district chiller for off-peak charging was €349 and €310 for the baseline and optimisation modes, respectively. The electrical tariff is shown in Table 2 (PNM 2011). This reduction equaled a total cost savings of about 11 %. The average cost per day was calculated to be €12.8 (baseline) and €11.46 (optimisation), and each mode had similar standard deviations. This suggests that the optimisation provides value on a daily basis as well as in the overall cost savings.

Table 2. Cost of electricity for the MEBldg.

June, July, August On-Peak Energy (€/kWh)	June, July, August Off-Peak Energy (€/kWh)	Other Months On-Peak Energy (€/kWh)	Other Months Off-Peak Energy (€/kWh)
0.061	0.025	0.048	0.025

The total and daily costs for the optimisation mode were lower than the baseline, but because the experiment was conducted over a limited timeframe of eight weeks, proper evaluation of the data requires normalisation of the daily costs with respect to the thermal load of the building, as shown in Figure 10 right. The figure also shows two regressions for optimised and baseline cases. The optimised regression line has a low R^2 value and the line has a slight upward trend as the load increased. The regression line for the baseline data reported an even lower R^2 value and an even lower slope. Although the regression lines for each data set did not result in a good fit, some hints can be inferred. For instance, at low building loads the optimisation offers a cost savings potential that could reach €3.75 per day. Also, as the thermal load increased, the regression lines trended towards each other and indicated that the opportunity to save money decreased as the thermal load increased.

Conclusion

This project has developed a framework for optimal scheduling software for buildings and microgrids using a SaaS paradigm, with the assistance of a third party energy services provider. The work extended previous MEBldg endeavours with the primary goal of increasing the robustness of the data transmission process. It also measured the cost savings potential between baseline operations and DER-CAM schedules. These experiments indicate that the reliability of the daily scheduling was less than 50 %, and yet, optimised schedules could still be implemented and beneficial because in part of the rolling 7-day outputs stored at the local facility. The second experiment determined that the DER-CAM schedules provided a substantial cost savings of 11 % for the MEBldg, in spite of low reliability for the transfer of schedules to the local facility.

The work also defined results of individual component performance, such as load forecasts and equipment operations. The performance and reliability results provide insight in where to make improvements and concentrate future work. For instance, the reliability of DER-CAM requires maintaining server to server connectivity through robust error checking and server maintenance. Other fertile areas for continued work are load forecast improvement, and equipment models that better match actual performance given numerous operational constraints. This could include increased feedback from building sensor data to monitor performance. One such particular interesting area is storage SOC. Controlling and tracking SOC could result in significantly better performance, and work is progressing on methods for achieving this.

Overall, schedule delivery and implementation was poor. In other words, the optimal schedules either provided or potentially provided (when its execution failed) are not being accurately delivered and executed. On the other hand, significant benefits were achieved, as described in the results section. This implies

that performance improvement is quite possible. Whether costly improvements are justified is quite a different question. Difficulties of implementation that derive from poor representation of devices and constraints on their operation, or these problems themselves, may not be easily resolved, and experience at any site may not be transferable. Since the software execution is potentially very cheap, the focus must be on capturing any attractive benefit from it, not on complex enhancements.

Glossary

AHU	air handling units
ASHRAE	American Society of Heating, Refrigerating, and Air-Conditioning Engineers
CERC-BEE	U.S.-China Clean Energy Research Center on Building Energy Efficiency
CHW	chilled water
DER	Distributed Energy Resources, includes local generation, storage, and controllable load
DER-CAM	DER Customer Adoption Model
DER-CIP	Distributed Energy Resource Customer Integration Provider
DES	District Energy System
EMCS	energy management and control systems
HVAC	heating, ventilation, and air conditioning
LBNL	Lawrence Berkeley National Laboratory
MEBldg	Mechanical Engineering Building
OSP	One Sun Plaza
SaaS	software-as-a-service
SOC	state of charge
UNM	University of New Mexico

References

- A. Mammoli, P. Vorobieff, H. Barsun, R. Burnett, D. Fisher. Energetic, economic and environmental performance of a solar-thermal-assisted hvac system. *Energy and Buildings*, 42 (9), 2010.
- A. Mammoli, M. Stadler, N. DeForest, H. Barsun, R. Burnett, and C. Marnay, Software-as-a-Service Optimised Scheduling of a Solar-Assisted HVAC System with Thermal Storage. MICROGEN III: The 3rd International Conference on Microgeneration and Related Technologies, Naples, Italy, 15–17 April 2013.
- Chris Marnay, Michael Stadler, Afzal Siddiqui, Nicholas DeForest, Jon Donadee, Prajesh Bhattacharya, and Judy Lai. “Applications of Optimal Building Energy System Selection and Operation”, *Journal of Power and Energy*, vol. 227 (1), 2013.
- Chris Marnay, and Judy Lai. “Serving Electricity and Heat Requirements Efficiently and with Appropriate Energy Quality via Microgrids,” *Electricity Journal* vol. 25 (9), 2012.

- Chris Marnay, Giri Venkataramanan, Michael Stadler, Afzal Siddiqui, Ryan Firestone, and Bala Chandran. "Optimal Technology Selection and Operation of Microgrids in Commercial Buildings", *IEEE Transactions on Power Systems*, 10.1109/TPWRS.2008.922654, vol. 23 (3), Aug 2008.
- David Steen, Michael Stadler, Gonalo Cardoso, Markus Groissb ock, Nicholas DeForest, and Chris Marnay, "Modeling of thermal storage systems in mixed integer linear program distributed energy resource models", *Applied Energy*, 2014.
- Douglas Montgomery, and George Runger. *Applied Statistics and Probability for Engineers*. John Wiley and Sons, Inc. 6th ed. 2014.
- <https://www.pnm.com/rates>. Rate 15B 2011.
- M. Ortiz, H. Barsun, H. He, P. Vorobieff, A. Mammoli. Modeling of a solar-assisted HVAC system with thermal storage. *Energy and Buildings*, 42 (4), 2010.
- Marcus Jones "Coupling TRNSYS and MATLAB for genetic algorithm optimisation in sustainable building design." Third German-Austrian IBPSA Conference Vienna University of Technology, 2010.
- Nicholas DeForest, Michael Stadler, Gonalo Cardoso, Tobias Brandt, and Sankar Narayanan. "Enabling broad adoption of distributed PV-storage systems via supervisory planning & control," ACEEE Summer Study on Energy Efficiency in Buildings 2014, Asilomar Conference Center Pacific Grove, CA, August 2014.
- P. Armenta, P. Vorobieff, A. Mammoli. Summer off-peak performance enhancement for rows of fixed solar thermal collectors using flat reflective surfaces. *Solar Energy*, 2011.
- S. Kiliccote, M. A. Piette, D. Hansen, Advanced Controls and Communications for Demand Response and Energy Efficiency in Commercial Buildings, 2nd Carnegie Mellon Conf. in Elec. Pwr. Syst. January 2006.
- US Dept. of Energy, Commercial Buildings Energy Consumption Survey, 2012.

Acknowledgements

This work was primarily funded by the U.S. Department of Energy Buildings Technology Office U.S.-China Clean Energy Research Center on Building Energy Efficiency (CERC-BEE) program <http://china.lbl.gov/research-projects/us-china-cerc-bee-program> The same technology will be demonstrated in one of 5 Chinese demonstration buildings with industrial partners, Johnson Controls Inc. and Xing Ye Solar, partially funded by China's Ministry of Housing and Urban Development. Additional funding was provided by PNM to UNM and LBNL under the Cloud-Based Energy Resources Scheduling project. Nicholas DeForest and Michael Stadler contributed to earlier work with the UNM MEBlg.