Lawrence Berkeley National Laboratory

LBL Publications

Title

Moore vs. Murphy: Tradeoffs between complexity and reliability in distributed energy system scheduling using software-as-a-service

Permalink

<https://escholarship.org/uc/item/4s73f3m7>

Authors

Dutton, S
Marnay, C
Feng, W
et al.

Publication Date

2019-03-15

DOI

10.1016/j.apenergy.2019.01.067

Peer reviewed

Moore vs. Murphy: Tradeoffs between complexity and reliability in distributed energy system scheduling using software-as-a-service

Spencer Dutton^{a,*}, Chris Marnay^a, Wei Feng^a, Matthew Robinson^b, Andrea Mammoli^b

^a China Energy Group, Lawrence Berkeley National Laboratory, 1 Cyclotron Rd, MS 90-1021, Berkeley, CA 94720, USA

^b Mechanical Engineering, University of New Mexico, MSC01-1150, Albuquerque, NM 87131, USA

ABSTRACT

Software-based optimization of building control strategies, including scheduling, has the potential to improve the performance of existing complex heating, ventilation, and air conditioning (HVAC), storage, and other systems—especially if temporally variable energy production, such as solar thermal or photovoltaics, is included. If reductions in energy bills can be achieved using optimized control strategies that take advantage of cost-saving opportunities, such as time-of-use pricing, the additional bill savings can cover further efficiency investment costs. As computer processing becomes cheaper over time (Moore’s Law), opportunities to perform complex control optimization become more abundant, and these can be performed remotely as software-as-a-service (SaaS). However, by “perfecting” our control strategies, we run an increased risk that when something unexpected happens (Murphy’s Law), the consequences of failure are greater. This study used simulation to explore the potential benefits of HVAC schedule optimization, delivery, and implementation using a SaaS paradigm, at various levels of complexity. Implementing optimal schedules in a model of an efficient building’s HVAC system, the study predicts energy cost savings of up to 10% compared to the naïve reference control strategy. Optimizing more system control variables increases the potential energy cost savings; however, these savings could be compromised by failures in communication inherent in delivering schedules via SaaS. The additional cost of energy resulting from the risk of increased demand charges generally increased with increased communication failure to a much larger extent than the risk of increased energy use charges. This work suggests that moderate improvements in performance, achieved at low cost by simple means, may be more effective than highly optimized schemes, which are more susceptible to failure due to their dependence on complex interactions between systems.

1. Introduction

1.1. Background

A common misconception associated with building energy use, especially electricity, is that reducing energy consumption through efficiency improvements is always the primary means of reducing energy costs; whereas, in reality, in a complex energy system, energy

consumption is one of multiple factors that drive the economics. Historically, and still predominantly, researchers assume that the investment funds needed to pay for efficiency improvements can be recouped by future energy expenditure savings; however, as buildings become more efficient, the energy conservation fruit hangs increasingly higher. Inevitably, the building reaches a break-even point where energy bill savings can no longer justify further investment to improve efficiency.

* Corresponding author.

E-mail address: smdutton@lbl.gov (S. Dutton).

One promising option to further reduce costs associated with energy use is to drive down the cost of implementation, which effectively shifts this breakeven point, allowing an improved return on investment. At the same time, non-linearities in billing, such as time-of-use (TOU), demand charge, or other complex tariffs, can be exploited to generate additional savings. In other words, energy management measures using price variations in tariffs provide additional opportunities to deliver cost savings in low energy use buildings in addition to measures that are purely efficiency motivated. Of course, the lower the energy use of a building, the smaller the potential energy efficiency bill savings; therefore, as buildings become very energy efficient savings captured in non-energy-saving ways become an ever bigger share of the potential cost reduction.

This thinking leads to a search for inexpensive cost-saving measures, such as behavioral changes or, as in the case of this work, software solutions that can improve the performance of existing hardware. The software-as-a-service (SaaS) model can serve to help here because its execution could be very inexpensive, effectively no more than executing a job on a server somewhere in the cloud, where the main cost consideration is the remote computation cost. As these costs fall over time, courtesy of Moore's Law, the opportunities to perform SaaS based optimization only increase. While setting up such a problem is a highly specialized task, executing it can be trivial or automated, and the only real challenges with exchange of results are reliability and cyber security. It could be reasonably argued that as systems become increasingly complex, they are likely to become less reliable as the number of potential failure points increases, leading to the popular wisdom that if anything can go wrong it will, i.e. Murphy's law. These two counter-acting forces lead to the motivation of this paper and its title. On the one hand, complex remote SaaS system optimization can lower energy bills with minimal investment and negligible operating cost, while on the other, the necessary multi-step data link will be error prone so any cost gains can easily evaporate.

Particularly difficult, however, is implementing the abstracted remote modeling results in real local legacy systems, which is where this work is focused. Crucially, all schedule optimizations are cost minimizing (i.e., not energy minimizing) making the electricity tariff one of the key inputs. The potential gains from optimization will be most attractive for complex building systems operating under complex tariffs. Recent trends indicate that building-sited energy resources, or distributed energy resources (DER), will become increasingly commonplace, supplementing traditional centralized electricity delivery infrastructure. Medium and large commercial facilities are prime candidates for this shift in how energy services are provided and used. When deployed appropriately, DER has the potential to offer benefits to the owner of the facility, in terms of local control, reduced costs, reduced emissions, and increased reliability and resilience. To maximize these benefits, DER operations should be optimized. Even for relatively simple systems though, optimization is far from trivial, typically requiring the solution of a mixed-integer constrained optimization problem, making SaaS an attractive means of delivering optimized schedules to commercial facilities. To evaluate its real-world effectiveness, researchers at the University of New Mexico (UNM) and the Lawrence Berkeley National Laboratory (LBNL) collaborated to implement interfaces between the Distributed Energy Resources Customer Adoption Model (DER-CAM), and the local building automation system (BAS) of the campus Mechanical Engineering building (MEB). DER-CAM is a mixed-integer linear program (MILP) developed over many years by LBNL. The DER-CAM server is the optimization engine used in operations mode; in other words, in the modeling applied here, the physical building energy system is a given to DER-CAM which only attempts to operate it in the cost minimizing way without any equipment changes.

The building's heating, ventilating, and air conditioning (HVAC) system, shown schematically in Fig. 1, incorporates cooling assisted by a 232 m² solar thermal array providing heat to a 70 kW_{th} absorption chiller, which also provides wintertime heating. A 30 m³ hot thermal

storage tank makes heat available later for both heating and absorption cooling, and a 350 m³ chilled water (CHW) storage is used to shift cooling electrical loads to off-peak periods. Data on the performance of various technologies (e.g., storage tank losses, pump power consumption) were obtained either by direct measurement or from the technical literature. A campus district energy system (DES) powered wholly by grid power supplies CHW and steam to all buildings on campus via heat exchangers. During the cooling season, mid-March to mid-October, the building can be cooled by its large thermal cold storage, the solar-powered absorption chiller, the campus DES CHW, or by a dynamic combination. The solar absorption system provides approximately 40% of the building cooling load on an annual basis. Cooling not provided by the absorption chiller is delivered by the storage tanks, and additional cooling can be supplied by the DES via a heat exchanger normally used to charge storage off-peak. This level of complexity, especially the need to optimize storage charge and discharge scheduling under a TOU tariff with a monthly ratchet demand charge, presents an ideal application for schedule optimization as performed by DER-CAM.

This study emerges from a demonstration to optimize the operational control of the MEB that was executed periodically over two years. The field approach employed a SaaS system that directly delivered cost-minimizing operational schedules to the building's BAS. In this approach, data were exchanged between a local control system and a remote optimization engine. In the summer of 2014, a demonstration of the effectiveness of DER-CAM to deliver energy cost savings at MEB was carried out, by comparing energy costs for alternating weeks in which DER schedules were either fixed (as would be implemented in a typical BAS) or optimized daily by DER-CAM [1]. Specifically, the schedules controlled the cold Thermal Energy Storage (TES) daily charge target and operating times for the absorption chiller. This demonstration effectively lowered energy use and generated energy bill savings up to 30% in the most favorable weeks with an average of 11–15% for a sequence of summer on-off test weeks. While the savings are potentially big, the results of this work led to an understanding of the many practical modes of failure of this approach and the likelihood that an individual optimized schedule would be successfully delivered and implemented. Results showed that, while models of building systems can be built successfully and remote optimization executed at very low cost, delivering the schedules electronically from the optimization server to the client system, and successfully executing them, can be a challenge making the resulting controls unreliable. Over a period of 24 months of periodically performing the SaaS-based trial, successful delivery and implementation of the optimized schedules occurred only around 50% of the time, corresponding to a failure rate of approximately 15 failures per month. One of the most significant findings was that the reliability of communications between the optimization server and the end user was a key determinant of success. Note that the benefits are easily sacrificed by small failures. Particularly, the demand charge for the month could be set by one high electricity load period (one hour in this case) caused perhaps by exhausted thermal storage resulting from failed delivery of the day's correct charging target. Perfect schedule delivery and excellent load control in every other hour of the month cannot reverse the damage.

Field experience led to this study, an evaluation of the real-world value of improved communications performance and operational detail under tightly controlled conditions. The evaluation is carried out by simulating the interaction of the SaaS system serving MEB and the operation of its subsystems. The mathematical treatment of the system requires consideration of both the physics of individual devices as well as the relationship between devices. For the purposes of this work, the physical system is modeled using a well-established commercial software (TRNSYS) that was developed specifically for tasks such as this. The specifics of the model components used for the simulations described in this paper are provided in Appendix A. The correct implementation and operation of the model was thoroughly checked by ensuring that model results were consistent. For example, the energy

Fig. 1. Schematic of the UNM MEB solar-assisted HVAC and energy storage system.

removed by the heat exchanger component X2 from the cold storage tanks can be calculated by integrating the product of flow rate, temperature difference and specific heat capacity over the charge period. This quantity was compared to the change in heat content in the tanks themselves, based on initial and final temperature profiles, in combination with specified heat loss through the tank walls.

Total energy costs are calculated as a function of the reliability of schedule delivery and schedule complexity. By simulating a large number of outcomes with stochastic communications success and varying schedule complexity, the nature of the trade-offs between improved performance and risks are explored. What makes these potential faults much more interesting is the asymmetry of consequences. In efficient buildings, energy costs are already low, so small percentage deviations in usage in either direction do not have a big effect on the energy bill; however, these deviations can have a major impact on demand charges, as noted above. The more tightly a highly efficient building is managed, the less important controlling energy becomes compared to controlling power, as noted above. One peak load slip-up can irrevocably change the monthly bill. In other words, the optimization benefits are significant but uncertain.

The main objective of this second stage of work is to build upon the original field study, to analytically investigate the failure energy and cost implications of the SaaS-based control optimization strategy, while exploring alternative strategies that optimize more of the HVAC system parameters. To remove uncontrollable parameters inherent in a purely experimental approach, including the inability to compare performance under identical conditions and the expenditure in time required, the study was performed using a high-fidelity TRNSYS model of the MEB rather than the building itself. Consequently, the repeat experiments can be simulated under identical conditions of weather and internal loads while exploring the parameter spaces of schedule complexity and schedule delivery reliability. While in the real system failure could result from multiple causes, for this work, failure is defined as the inability of the IT framework to deliver a rolling 7-day-ahead schedule on any particular day. The failure rate for a complex system can be derived using a steady-state analysis based approach as presented by Koepfel et al. [2]. In the work presented in this paper, the focus of our analysis was not on the measured failure rate of the real world system installed in the MEB, which, due to the research nature of the technology experienced significantly higher failure rates that would be tolerated in a mature technology, but rather on the simulated outcomes of a broad range of hypothetical failure rates. Although motivated by real field experience, and the modeling is based on a real building and a real SaaS set-up, this work is a purely analytic exercise.

The experimental framework is presented first in the methodology section, outlining the model of the building, the method by which optimized schedules are produced, and how uncertainty is introduced. The results section quantitatively explores the potential savings that can be achieved by careful scheduling of the equipment, such as could be done by an experienced building manager. Optimized schedules are then applied, comparing the cost benefits to what is achievable without optimization under various levels of uncertainty in schedule delivery and load forecast accuracy. Finally, the implications of the results are discussed as well as additional sources of reliability issues that should be addressed in further work.

1.2. Literature and contribution

This work draws on three distinct areas of literature, building (especially DER) operational optimization, building energy modeling, and finally the analytic question at the heart of this work, the financial risks associated with unreliability.

The DER-CAM decision support tool [3] was used as the operational optimization engine for this work. Built on the GAMS platform [4], this tool has been developed over many years by LBNL, and it is a leading example of MILP DER optimization. Its objective is to minimize the annual costs for providing energy services to the modeled site, including utility electricity and natural gas purchases, amortized capital and maintenance costs for DER investments, subject to a variety of constraints [3]. Very similar mathematics can be used to either optimize the operation of an existing system, i.e. over the economic short-run, or to optimize both system operation and investment over time, i.e. the economic long-run. Various other software tools are available to meet both these requirements. HOMER [5] for example, is a widely used DER planning package. However, this study only involves operational optimization of a fixed system over a moving 7-day horizon, which is ideally suited to using DER-CAM.

Earlier work [6–8] on U.S. adoption of distributed energy resources showed that the technology and economics of electricity production and delivery are tilting away from huge scale centralized provision and towards local generation close to the point of end use. This shift suggested a quite different approach to planning and operating electricity generation and delivery is needed. Rather than developing the traditional remote centralized power system more-or-less independently of other energy systems, the optimization of electricity generation and use in close proximity should be brought into the systemic framework of building energy modeling and integrated with optimization of other

energy systems. Marnay et al. [3] find that this shift is particularly noticeable in organization of commercial building scale systems; where a key early integration is building electrical and heat systems [9]. The scale of electrical and thermal system integration and its optimization extends beyond individual buildings to the campus and community scale in the form of district heating and related networks, which are widespread and have a long history, and more recently similar methods are being applied to low-carbon communities [10–13]. To date, energy storage is rare in traditional utility-scale systems. Work by Stadler et al. [14,15], DeForest et al. [16] and Steen et al. [17], show that when considering DER, incorporating energy storage provides additional utility compared to more stable centralized systems, but incorporating it presents a new set of optimization challenges. This is particularly the case where energy storage is not just limited to electrical systems, i.e. batteries, but can also encompass thermal, including in building mass and other forms. Note that in this work, both these aspects are in play; the UNM MEB has solar energy collection as well as electrically driven cooling whose operation must be co-optimized with electricity use, and optimally operating its huge thermal storage system is central to a cost minimizing solution. The MEB is a 7000-m² *living laboratory* for advanced building technologies. The building, commissioned in 1980, received a thorough modernization between 2006 and 2010. The MEB was modeled using TRNSYS, as described in several earlier works by Ortiz et al. [18] and Mammoli et al. [19,20]. TRNSYS allows the construction of complex models of buildings and energy system components therein, using a graphical user interface in which components can be connected via physical (energy and matter) and information links, as would be done in a real system. Each component is described by relevant differential equations, that are solved numerically, often using spatial discretizations, which together with material parameters and component technical specifications can be set by the user. More specific details on equations that describe the model components, such as the heat storage tanks or the solar array, can be found in the literature (e.g. [21]).

Recently, work by Cardoso et al. [22], and Marnay et al. [23], find that, other aspects of localized power supply have received considerable attention, notably the reliability and resilience of DER. Work by Stadler et al. [24], Jin et al. [12], Nemati et al. [25], Li et al. [26], Quashie et al. [27] and Wang et al. [28], also explore the integration of demand response into a single optimization, with new methods being applied to the problem, including multi-level optimization.

Previous studies have found that uncertainty can lead to poorer than expected performance in optimized scheduling. For example, Gao et al. [29] find that load prediction uncertainty can reduce performance of a demand response program (DRP) in terms of cost savings by a factor of two. When robust control using a Monte Carlo approach is used, the DRP can still produce a saving close to the ideal savings obtained with a perfect forecast. Wang et al. [30] find that uncertainty in model accuracy, model input parameters, weather and operational practices can all affect the energy consumption of the building. Using a Monte Carlo approach, they conclude that uncertainty in energy consumption follows a log-normal distribution, with a range that can dwarf the effect of design features.

The approach taken here structurally resembles the information-gap decision theory (IGDT) paradigm, in which a model of a system has some unknown parameters. The uncertainty in parameters is modeled by nested subsets around an estimate, and finally the ability of the system to satisfy performance requirements is measured. The size of the subsets grows with increasing uncertainty. Robustness of a decision is a function of the size of the subsets that still produce acceptable performance, and it is independent, in a sense, of the distribution of parameters. A recent application of IGDT is found in the work of Ahmadi et al. [31], in which the operating strategy of generating units (in the form of a day-ahead unit commitment) within a distribution system with battery storage is adjusted based on the acceptable level of risk against load demand uncertainty. A similar case is considered by

Kazemi et al. [32], using robust optimization, in the form of a min-max problem. The case of day-ahead schedules for an apartment building with battery and thermal storage, CHP, and solar thermal collectors is treated by Najafi-Ghalelou et al. [33,34] who use IGDT to evaluate robustness against market price variation. Use of IGDT produces substantially better robustness to price increases. Similarly, the ability to take advantage of decreases in energy costs improves using IGDT. Mehdizadeh et al. [35] propose a short-term generation scheduling for a grid-connected microgrid in a day-ahead market in the presence of a demand-response program (DRP) under uncertainty. IGDT-based robustness and opportunity functions are proposed to obtain an optimal bidding strategy. Under these conditions, DRP can lower costs by almost 5%. While there is ample evidence that uncertainty can have a substantial effect on performance, the interaction between uncertainty and complexity is not addressed in the literature. The present work is intended to provide some insight into this phenomenon, and contributes to this discussion by addressing both the reliability implications of relying on remote operational optimization, as well as the trade-off between operational cost saving and exposure to failure. The present study also extends analysis of SaaS systems, for which DER-CAM has been previously used [1,36]. Indeed, the poor reliability and its consequences described by Jones et al. [1] are central to the motivation for this paper. Other efforts to optimize local power systems using cloud resources are also underway [37,38].

An active area of current research extends the bounds of integration to include electric vehicles, which will have an impact on both the wider power system and the buildings where they are interconnected, both as loads and potential resources. Optimization of this system that overlaps the buildings and transportation sector opens a new horizon, and given the importance of transportation decarbonization, a critical and enticing one [39].

This work furthers the state-of-the art by using simulation tests to address the Moore versus Murphy tension observed in an earlier field demonstration. As such it explores and quantifies a potential limitation of SaaS, within the context of the problem at hand. Matching the limited analytic MEB representation in DER-CAM, with the detailed model of the physical system in TRNSYS, requires thoughtful and technically accurate implementations in both environments. Finally, the major contribution addresses the implications of uncertainty in IT performance related to the delivery of optimized HVAC operating schedules.

2. Methodology

First, the nature and outcome of the original field study that motivated this experiment are summarized, including the methodology used to generate optimized equipment operations schedules and the physical equipment.

Using a TRNSYS model of the MEB HVAC components that replicates the original physical setup, implications of alternative control scenarios in simulation were examined. The hypothesis tested is: as the number of different components of the system that rely on optimization increases, so do the consequences of failures in that optimized control strategy (in terms of energy cost). The hypothesis also anticipates that increased optimization leads to increased savings when everything works well.

To test this hypothesis, three different control scenarios were developed with an increasing number of factors being optimized, representing increased levels of complexity. In the simplest scenario, **S0**, only the cold TES target charge level is considered. In the second scenario, **S1**, the absorption chiller operation schedule is also considered, while in the third scenario, **S2**, the time at which the cold TES charge begins is additionally included in the schedule implementation. The model is applied first, showing that, at least in the deterministic case of perfect load forecast and perfect schedule implementation, optimization at increased levels of complexity does produce lower costs. A Monte Carlo approach is used next to assess the cost sensitivity of these

three scenarios to failures, to successful implementation of the optimal control schedule, and to inaccuracies in the forecast load.

2.1. Field study apparatus

In the original UNM-LBNL field study, a model of the MEB envelope in TRNSYS [1] was developed. This model predicted seven-day-ahead forecasts of building thermal loads based on forecasted weather data obtained from the National Weather Service. These load forecasts, along with environmental data, the mechanical system definition, system parameters, and energy tariffs are used to generate GAMS configuration files. These files are transferred each day to the DER-CAM server at LBNL, which optimized operational schedules for a rolling one-week forecast period. Finally, the schedules were retrieved from the LBNL server and inserted into the BAS that controls the operation of the MEB's HVAC system using an open database connectivity (ODBC) platform provided by the BAS vendor. While this arrangement may seem unnecessarily complex, involving several potentially error-prone interfaces, it reflects possible real-world implementations. Typically, DER are being deployed into a building environment with legacy systems installed that make achieving interoperability notoriously difficult. A more detailed description of the SaaS architecture used to deliver optimized operational schedules from the remote LBNL server to MEB's BAS is given in Jones et al. [1].

2.2. Detailed simulation method

For the present simulation stage of this study, the same SaaS backend is used to generate optimal system operation schedules; however, instead of using them to control the real building, they are used as inputs to a calibrated model of the building's HVAC system. The model includes detailed descriptions of the stratified thermal energy storage (both hot and cold), solar array performance, air handler performance, and pumping electrical loads. This HVAC model was also used in TRNSYS as the basis for the simulation study. Details of the HVAC system model used in this work are presented in Appendix A.

The simulation was performed in two stages. First, based on prior field study experience, some initial baselining studies were developed to confirm correct operation of the model and to confirm that, in principle, schedule optimization does deliver savings. General changes to operational strategy were considered that intuitively should deliver cost savings, changes that are simple enough to be implemented without the aid of the DER-CAM optimization engine. This stage of the work is described in Section 2.3 Baseline operation.

Second, Section 2.4 Optimization and uncertainty modeling introduces the optimized schedules generated by DER-CAM. These schedules and their corresponding cooling load profiles were applied in the model, at first perfectly (i.e., for the case where the actual building load is as the forecast, and schedules are always delivered), and then with various levels of synthetic error added. The error corresponds to load forecast inaccuracy and to schedule delivery failures. The error in load forecasts was based on previous experience with building operations [1], while the probability of successful schedule delivery is one of the experimental control variables varying between the previously observed value of 0.5 and the ideal unity case.

2.3. Baseline operation

Before evaluating the effect of optimized schedules on operating costs, a set of baseline scenarios is considered that could be implemented manually by an experienced systems operator. The reference simulation scenario, **Ref_baseline**, is based on the default operating mode of the MEB prior to the introduction of DER-CAM scheduling. Then, a set of strategies that represent incremental changes to this reference scenario, one control variable at a time, is developed. The objective of this incremental approach is to verify whether or not, and at

what level, computational optimization of these control variables can be expected to deliver efficiency improvements.

The next hypothesis, based on prior experience with chilled water storage systems supplemented by absorption cooling, is that three main sources of inefficiency lead to increased cost. First, over charging the cold TES to full capacity each night results in wasted cooling capacity at the end of the day. Although the majority of this leftover capacity is, in principle, available for use the next day, it is degraded by standby losses through the tank walls and diffusion across the thermocline, which would not occur if the TES were depleted daily. Optimizing the state of charge based on the following day's forecasted load presents an opportunity for savings. The second source of cost inefficiency considered is the use of the chiller during periods of peak electricity pricing. Even though the chiller's primary source of energy is heat, there are ancillary electric pumping costs; therefore, shifting a portion of the chiller schedule to the off-peak period can reduce cost, while it may or may not lower demand cost. The third source of inefficiency considered is the scenario where excessively early onset of the TES charging results in cold storage degraded by conduction through at the tank boundary, particularly the edges. It should be noted that the un-insulated tank walls (0.3 m thick concrete) are exposed to fresh air and exhaust chases, with an average ambient temperature of 25 °C, so standby losses can be considerable.

The reference **Ref_baseline** control strategy was to fully charge the cold TES during the off-peak period beginning at 8 pm and to operate the absorption chiller on-peak. The charging process begins a short time after the onset of the off-peak electricity rate, at 20:00 h, and stops when the temperature at the top of the tank falls below 12 °C. Operation of the absorption chiller begins as soon as the solar array has provided sufficient stored heat in the hot water storage tank to allow the system to "ride through" cloud-driven intermittency for at least one hour. The onset of chiller operation therefore depends on weather conditions, primarily sky cover. This reference model was used to assess *baseline* savings potential of the scheduling, as would be obtainable by expert programming of the BAS, with the expectation that an actual optimization of the schedules would deliver greater savings.

To represent the intended effect of reducing carry-over charge, the scheduled state of charge was lowered to a flat 60% of full capacity (**Ref_const_60%**). As a second step, it was further reduced to 0% at weekends, maintaining a 60% target charge during weekdays (**Base_S0**) throughout the simulation period. Note that the charging control does not know the current state of charge and attempts to deliver the equivalent thermal energy to a 60% charge; however, if the tank reaches maximum capacity because of carryover from the previous day, charging ends prematurely.

To account for improved chiller operation, the chiller was allowed to turn on in the early morning hours from midnight until 6 am, reducing on-peak activation time. This results in several hours of chiller operation being shifted to the off-peak early morning hours. This shift brings the chiller operation closer to the point of chilled water use the following day, reducing thermal losses. This strategy, **Base_S1**, included all the changes made in the **Base_S0** strategy, with the additional changes to chiller use scheduling.

In the third strategy, **Base_S2**, the evening tank charging is shifted from starting at 8 pm to starting at 12 midnight, until the onset of higher daytime energy rates the following morning at 8 am. This strategy includes the changes of **Base_S1**, but shifts the charging to a period just before cooling is required the next day, reducing storage losses. Moreover, the chiller operation is moved to the period from 8 pm to midnight.

The final strategy, **BS_FailCase**, is intended to demonstrate the effect of incorrect optimization of one of the key optimization variables, the tank state of charge. In this strategy the **Base_S2** strategy was changed such that the maximum scheduled charge was 50% of full capacity. Because this level of charging is insufficient to meet the cooling load of the building, auxiliary cooling from the DES is activated on-peak, resulting in additional demand charges.

2.4. Optimization and uncertainty modeling

As with Baseline operation, the TRNSYS model of the buildings HVAC system was used to explore scenarios with progressively more aggressive optimization, this time using the operation schedules generated by DER-CAM. As was the case in the field study, the electricity tariff at any given time of day is a key factor driving this cost-based optimization. TOU electricity rates implemented in the model were based on commercial rates from Public Service Company of New Mexico (PNM), specifically rate 15B for large services and public universities. Energy use charges were modeled using \$0.0821025 per kWh on-peak (8 am–8 pm), \$0.0327765 per kWh of off-peak, and a single monthly ratcheted demand charge of \$9.56 per kW on-peak. The generated schedules are optimized to minimize the total energy bill including charges.

Three scenarios were considered, representing increasing levels of complexity and shown here by the number of factors that are being optimized. In scenario **S0**, the schedules determine how much to charge the cold thermal storage each night, with the aim of minimizing over charging. Scenario **S1** adds control of the absorption chiller start and stop times; scenario **S2** adds optimization of the cold TES charge start time.

Scenarios **S0**, **S1**, and **S2** assume that the schedules were successfully implemented 100% of the time. Perfectly implemented scenarios are referred to as **S0 P1.0**, **S1 P1.0**, and **S2 P1.0**. At this step the TRNSYS model is calibrated to adjust for the inherent differences between the relatively simple DER-CAM model and the more realistic non-linear TRNSYS model. The tank charge schedule is a fraction of the total charge. To convert to an actual charge capacity in kWh, the charge fraction should be multiplied by a scaling factor 4095, calculated using the volume of the tank and the specific heat capacity of the water. To calibrate the TRNSYS model, this charge multiplication factor was adjusted up until the TRNSYS model no longer experienced demand charges. For the perfect (P1.0) scenarios **S0**, **S1**, and **S2**, and the uncertainty modeling that follows, multiplication factors of 5250, 5250, and 4965 are used, respectively. These multiplication factor calibrations represent a step that would be reasonable to perform in the real world, akin to commissioning a conventional HVAC system.

A Monte Carlo approach was then used to introduce synthetic errors into both the operation schedules and the forecast load file. One hundred unique 31-day operation schedules were generated for the three different levels of optimization complexity. Each day of the 31-day operating schedules were built up from the 31 days worth of seven-day-ahead forecasts. On each simulated day, a certain probability of failure in schedule delivery was introduced that would require the use of an older schedule. Typically, in reliability engineering, a systems failure rate is defined as the number of failures within a period of time as found in Hale [40]. In this analytical analysis, failures are defined as a failure to successfully deliver a viable operating schedule, where the probability of this succeeding has a predefined probability. Three probabilities were considered where, on any given day, the most up-to-date schedule was delivered correctly with success probabilities of 0.5, 0.7, or 0.9, which, given a scheduled update of one per day, corresponds to failure rates of approximately 15, 9, and 3 failures per month. For clarity, these failure rates are referred to in the remainder of this work by the probability of success of individual schedule transfers (0.5, 0.7, or 0.9). An illustrative example of one specific realization of a sequence of 60 daily schedules is shown in Fig. 2. The black “schedule transferred” line at the bottom of the plot shows the schedule actually used. The seven-day schedule delivered each day is represented by the stack of dots above each day. Days where the blue “one-day ahead” schedule was used indicate that a new schedule had been successfully delivered and implemented the previous evening. On days when delivery of the schedule was interrupted or failed, a previously delivered schedule was used (transferred from the previous day), as indicated by the various colored dots. Dots in the black schedule with color at the red end of the

spectrum are older and thus less accurate. Examples of multiple days of concurrent failures can be seen at days 8–9, 23–24, 28–29, 31–34, and 56–59. For example, for days 56–59, new schedules were not delivered for four days, meaning that schedules up to five days old are used.

Theoretically, failure to deliver on any one day has a small effect because a good substitute back-up is available from the previous successful delivery. Indeed, several back-ups of diminishing accuracy are available at all times. The consequence of a failed delivery, then, depends firstly on the accuracy of the best available substitute, which in turn depends on the accuracy of weather forecasts and other inputs. Secondly, failures might be considered random, and indeed they are in this work, but real-world failures are likely to be clustered around events, such as server failure. Failure patterns might have a significant influence on implications.

For each combination of three levels of complexity and three levels of risk of failure, 100 schedules were generated, resulting in 900 sets of schedules for the cold TES charge level, operation time of the absorption chiller, and TES charge start time.

To represent uncertainty in the forecasted building loads (i.e., the fact that the forecasted building load is different from what is actually experienced by the building), a pseudo random perturbation was added. This was done by adding a random walk to the forecasted building load for each day, with a mean of zero over large numbers of perturbations. Each of the 100 schedules at a complexity/success probability combination was associated with the corresponding “actual” building thermal load.

By simulating all of the 900 schedules using the TRNSYS model of the UNM building HVAC system, the performance of the system was investigated when following each of these stochastically generated schedules and loads, resulting in a probabilistic distribution of likely total energy bill. Results of this analysis were used to explore the probabilistic distribution of the energy-related costs that arise from single or multiple consecutive lost schedule events.

Table 1 summarizes of all the scenarios for both the *Baseline Operation* and for the *Optimization and uncertainty modeling* studies. Full simulations were performed for a total of 31 days to represent a full monthly billing cycle.

3. Results

3.1. Baseline

A summary of the results from the base-lining strategies is shown in Table 2, comparing the total energy use of the HVAC system, of the DES heat exchanger (HX) supplemental cooling, of the absorption cooling pumps, and associated costs. For the case of the **Ref_baseline**, in which the cold TES is fully charged every night during the off-peak period, the monetary cost for July is \$455 in energy use charges and \$53 in demand charges (based on electricity consumption of 12,316 kWh of which 1578 kWh resulted from on-peak use of the absorption chiller). Crucially, for the full period of operation, zero use of on-peak supplemental cooling from the DES was observed, because there was always an excess of cooling capacity in the cold TES. Reducing the scheduled state of charge from a constant 100% to 60% (**Ref_const_60%**) had no measurable impact on costs because the excess tank capacity was always available with significant carry over to the next day, meaning that the TES reached the maximum state of charge every night without needing the pre-defined 60% charge. Additionally, the low level of demand at the weekends acted to “top-up” the tanks.

For the **S0_base** strategy, the introduction of no charging on Friday and Saturday evenings resulted in reduced off-peak charging of the cold TES, reducing total energy-related costs by 7%. Also, with the **S0_base** strategy (as is the case with **S1_base** and **S2_base** below), supplemental cooling from the DES takes place on the weekends but does not incur demand charges because of the tariff design (weekends are considered off-peak, with no demand charges).

Fig. 2. Example schedule realization over 60 days for a 0.7 probability of success.

With the **S1_base** strategy, the cold TES charge process begins at 8 pm, four hours before the absorption chiller operation is scheduled to start. The chilled water tank reached maximum charge state before the absorption chiller is activated at midnight. The cooling delivered by the absorption chiller further reduced the temperature of the TES (i.e., the TES was overcharged) at much lower efficiency. While strategy **S1_base** was successful at reducing the costs associated with daytime use of the absorption chiller, the strategy failed to reduce overall energy costs as the reduction in supplemental cooling provided by the absorption chiller was made up by other means. Cost savings were obtained using strategy **S2_base**, in which the chiller operates partially off-peak, between 8PM and midnight, while TES charging occurs between midnight and 6am. The **S2_base** strategy results in a 9% decrease in the monthly bill over the reference case. For the scenario with an overly constrained tank charge (**S_FailCase**), the electrical power use is reduced compared to the **S2_base** strategy because of the reduced scheduled charge. Reduced charging has the effect of reducing electricity use charges; however, these are more than canceled out by the increase in demand charges as a result of *on-peak* operation of the supplemental cooling from the DES, resulting in a 1% increase in costs.

3.2. Modeling uncertainty in optimization

To evaluate the combined effect of optimized schedules and uncertainty in schedule delivery, first, the case where the load is exactly as forecast and there are no failures in schedule delivery was simulated. The operation of the TES, absorption chiller, and supplemental cooling is shown in Fig. 3. TES charging occurs during the night and ends before the cooling load begins each day. The absorption chiller operates partly on-peak and partly off-peak, as a result of the limited hot TES capacity. The demand charge is a result of the first occurrence of chiller operation on-peak. Cumulative energy costs rise steeply during cold TES charging (high power but low energy charge during off-peak hours) and less steeply during daytime absorption chiller operation (low power but high energy charges).

The cost of uncertainty can be observed by contrasting the operation and costs plot for the same period, but with uncertainty in schedule

delivery and in load forecasting accuracy. Operations and costs were higher energy and demand charges occurred as a result of uncertainty as shown in Fig. 4. In this case, higher demand charges result primarily from on-peak operation of the DES supplemental cooling as a result of insufficient cold TES charging (due to either a missed schedule or a poor load forecast). It is easy to see that demand charges can soon dominate costs. This is because even small errors in scheduling or load forecasting produce high cost increases due to the ratchet mechanism inherent in TOU tariffs with a demand charge.

As the operating schedules in the two scenarios are identical, demand charges were triggered in this case due to cooling load under-prediction, as shown in Fig. 5.

The effect of the imperfect load forecast is presented in Table 3. Energy use and energy cost results for the various optimization scenarios with optimized schedules including perfect load prediction appear alongside the Monte Carlo runs that include uncertainty.

In the fully deterministic case, the results show a significant (4%) reduction in total energy costs from the reference to strategy S0. Electrical energy use was reduced somewhat but all at off-peak rates. Strategy S1 provided a further 1% cost reduction by shifting the operation of the absorption chiller to off-peak periods. Finally, in the S2 scenario energy costs are reduced by an additional 5% through reduced standby thermal losses. When uncertainty in the load is introduced, substantial variation in the demand charge cost is observed. While the average is essentially equivalent for all complexity levels, the scatter is highest at the lowest complexity value and essentially equivalent for the higher complexity.

To better understand the statistical consequences of uncertainty in both the load forecast and schedule delivery reliability, energy and demand costs for each realization of a schedule delivery and load forecast are plotted against each other in Fig. 6 for complexity levels S0, S1, and S2, respectively. The plots show a number of interesting features. First, the optimized schedule with perfect load forecast and no schedule transfer failure always outperforms the baseline cost with daily full TES charge, by virtue of lower energy cost and equal demand cost. When load uncertainty is introduced while retaining perfect schedule delivery (the P1.0 cases), small changes in energy costs are

Table 1
Complete set of simulation scenarios and operating parameters.

	Scenario	Operating parameters
Base-lining scenarios	Ref_baseline	Full tank charge each night, absorption chiller available 24-7 (AC24-7)
	Base_Ref60	Tank top-up charge max 60% (TMC60), AC24-7
	Base_S0	No charge at weekends (NCW), TMC60 on weekdays, ab. AC24-7
	Base_S1	NCW, TMC60 on weekdays, chiller operation reduced peak hours (COH)
	Base_S2	NCW, TMC60 on weekdays, COH
	Base_Fail	
Complexity scenarios	S0	Tank charge max scheduled by DER-CAM (TCMD)
	S1	TCMD, Chiller operating schedule scheduled by DER-CAM (COD)
	S2	Tank charge time and max, and chiller schedule by DER-CAM
Uncertainty scenarios	S0 1.0 – P0.5	As per S0, with uncertainty in load and schedules
	S1 1.0 – P0.5	As per S1, with uncertainty in load and schedules
	S2 1.0 – P0.5	As per S2, with uncertainty in load and schedules

Table 2
Energy consumption and costs. Note that all energy values are in electric equivalent.

	Total electricity use (kWh)	HX supplemental (kWh)	Absorber pump energy (kWh)	Demand cost (\$)	Energy cost (\$)	Total cost (\$)	Reduction in total costs vs. reference
Ref_baseline	12,316	0	1578	53	455	508	
Ref_const_60%	12,316	0	1578	53	455	508	0%
S0_base	10,913	381	1578	53	422	475	7%
S1_base	11,061	391	1538	53	424	476	6%
S2_base	10,722	357	1538	53	411	464	9%
S_FailCase	10,398	428	1538	110	404	514	−1%

observed associated with often substantial increases in demand charges. Overall there is a trend of increasing demand cost with decreasing energy cost, but the reduction in energy cost is generally smaller than the corresponding demand cost. As the probability of timely schedule delivery is reduced, there is a general trend of reduced energy cost (i.e., the TES is not charged as much as would be needed) accompanied by often significant demand cost increases. On the surface, this behavior is puzzling: why should there be a reduction in TES charge when a schedule is missed? This research shows that each weekly schedule is based on the assumption that the TES is initially depleted. However, it is possible that at the end of the first day there is residual charge, which is considered in the optimized day 2 schedule but not in the new day 1 schedule. As the probability of timely schedule delivery is reduced further, the energy cost cloud compresses from the right against a lower bound of approximately \$370 for the month of July, while the demand cost cloud is fairly spread out between a rarely achieved minimum of \$50 and maximum of \$350. Occasionally demand costs exceed \$400. Interestingly, the realizations with lower probability of schedule success (P0.5) result in higher demand charges when the complexity is lower: for the S0 strategy, a maximum demand cost of approximately \$350 was reached several times, while for the S2 strategy the demand cost never exceeded \$300. An explanation for this is that, for the S2 strategy, the TES is charged at the optimal time to minimize standby losses, so that there is generally a larger energy buffer with the S2 strategy to counteract unexpected load and reduce the chance of large demand charges.

The distributions of total energy-related costs for the three levels of complexity can be represented using a box and whiskers plot, Fig. 7, showing median, upper, and lower inter-quartile and maximum and minimum total cost. The plot shows many interesting features. There is a clear effect of the probability of success of schedule delivery: the median energy cost decreases with increasing probability of success for all levels of complexity. The size of the inter-quartile box also decreases with increased probability of schedule delivery success, and becomes more skewed towards higher cost, as a result of the demand ratchet mechanism. Interestingly, for the lowest level of complexity, the size of

the whiskers increases with improved probability of schedule delivery success, which is counter-intuitive, but is likely, associated with uncertainty in the forecasting. In this case, using lower quality (older) schedules sometimes would produce overcharging of the cold TES, offsetting the under-predicted load.

Higher levels of complexity result in lower median cost at the same schedule delivery success probability. Also, at the higher level of complexity the size of the whiskers is lower, meaning that there are fewer outliers. One caveat, however, is that this work only considered one point of failure, the schedule delivery, which does not depend on complexity. In reality, the implementation of higher complexity measures would result in higher probability of failure.

By calculating the mean total energy cost of each of S0, S1, and S2 scenario distributions the relationship can be distilled between complexity, risk, and cost down to a more simple representation. Fig. 8 shows the relationship between complexity, indicated by the scenarios S0–S2, and uncertainty, given as the probability of success. Each point represents the averaged result from 100 runs. For the S0 scenario, with a limited daily risk of failure (P0.9), average costs exceeded the reference. When assuming perfect schedules and load prediction as shown in Table 2, the S0 scenario lowers costs compared to the reference. The introduction of these two categories of uncertainty result in an average increase in costs. The S1 scenario results in significant savings when the schedules were implemented perfectly; however, these rapidly become losses compared to the reference condition when failure rates increase. The S2 scenario showed the least sensitivity to increased risk of failure; energy related cost savings were observed when failure rates were 50%.

4. Discussion and conclusions

Results from the baseline study and models first showed that the model behaved as expected and that impressive savings up to 9% are possible using minimal control changes based on intuition; however, expecting the right intuition may be overly optimistic in practice, and will generally not be achievable by the typical building manager. Not having the right intuition could backfire, even under the best of

Fig. 3. Operation of the TES, absorption chiller, and DES supplemental cooling (left scale) and energy and demand charges as a function of time (right scale) for the case of perfect optimized schedule delivery and perfect load forecast.

Fig. 4. Operation of the TES, absorption chiller, and DES supplemental cooling (left scale) and energy and demand charges as a function of time (right scale), for a case of imperfect schedule delivery and imperfect load forecasting.

circumstances. On the other hand, applying optimization-based control strategies that vary the prescribed operating behavior hour by hour could deliver savings of up to 10% if forecasts are accurate enough and if system reliability is high. Unfortunately, in reality this is also seldom the case, and an aggressive optimization could produce unexpected negative consequences.

To quantify risks of excessively aggressive optimization, this study looked at implications of inaccurate load forecasts and of unreliable communications leading to missed schedule deliveries, in turn resulting in the application of less accurate schedules. While sometimes less accurate schedules can provide small savings from incomplete TES charges, these savings can be overwhelmed by high demand charges. In fact, one of the primary results this work shows that the cost risk of incurring extra demand charges is much more serious than the risk of consuming slightly more energy than needed.

Nonetheless, this work also produced a counterintuitive outcome: for the MEB case, aggressive optimization, which involves the control of several system parameters, is not necessarily more risky than a simpler optimization, and in fact can deliver more reliable savings.

The lowest level of complexity, in which only the total charge level of the TES is optimized, produces a moderate energy saving in the best of circumstances. However, when uncertainty in the load forecast and schedule delivery is accounted for, the expectation is that costs will be higher than for the base case over the long term, where neither intuition or optimization are used to reduce costs. Furthermore, when the probability of missed schedules is high, overall cost increases on average by 5%.

With a moderate optimization level (TES charge and absorption chiller schedule), expected costs are lower than the base cases if

schedule delivery is perfect, but rise above the baseline for higher risk of missed delivery. Surprisingly, only the highest level of complexity, involving TES charges, chiller schedules, and TES charge time delivers cost savings in the long run. This counterintuitive result was found to be a consequence of the higher “buffer” level of energy that can be relied upon to prevent demand charges.

A major observation from this study is that high costs result from the infrequent occurrence of demand charges for a very short period of time. The implication is that, to take full advantage of the results of optimization, some of the constraints could be relaxed. For example, thermal comfort for short and infrequent periods of time could be tolerated to ensure that energy savings are maintained. Alternatively, a small buffer could be implemented, for example by charging the TES a small amount more than needed to meet the next day’s load.

Another way of looking at these conclusions is that stochastic optimization is the key to reliable cost savings. However, stochastic optimization is more computationally expensive, and more challenging to set up. Perhaps simple rules or corrective actions based on deterministic optimization are a better practical option.

It is also shown that some feedback in the system may provide useful information for the optimization. For example, the assumption that the TES is completely discharged at the end of each day can lead to overcharging, which accumulates and eventually compromises savings from reduction of thermal losses. On the other hand, undercharging may result in excessive demand charge costs. The solution could be to use a known value of the state of charge of the TES at the end of the day, but in practice this requires additional instrumentation and associated hardware costs.

Finally, improving the reliability of IT systems should be possible

Fig. 5. Percentage difference in modeled air loads using a perfect and imperfect load forecast.

Table 3
Energy consumption and costs.

	P1.0 with perfect load prediction					P1.0 with load uncertainty [mean (min\max)]		
	Chiller energy kWh	Demand charges \$	Energy cost \$	Total cost \$	%Δ in total costs vs ref.	Demand charges \$	Energy cost \$	Total cost \$
Reference	1.58E+03	53	455	508	–			
S0	1.58E+03	53	435	487	4.0%	85(53\421)	437(432\441)	521(484\857)
S1	1.47E+03	53	430	483	5.0%	76(53\231)	432(427\436)	508(480\662)
S2	1.47E+03	53	404	457	10.0%	77(53\266)	406(401\409)	483(454\673)

and could be the least expensive option. Improving the reliability of load forecasts, for example using machine learning methods, could also be a low-cost option, maybe even lower than deterministic energy modeling of the facility. Higher complexity is not as dangerous as it would appear on the surface; in fact, improving the performance of a system from a variety of directions may make the system more robust. Finally, relaxing the constraints on thermal comfort (allowing room temperatures to rise above commonly accepted maximums), might be a better option than their uncompromising enforcement.

Although not addressed in this work, tightly controlled systems can potentially participate in other markets (e.g., for demand response or ancillary services, offering other promising but complex-to-manage opportunities).

Of course, this work considers one specific building and only scratches the surface of real world complexity. For example, only the probability of failure in schedule delivery is considered. The probability of failure in the implementation of individual measures, which would result in a higher probability of something going wrong with increasing complexity, is not considered. The measure of complexity utilized here, namely the number of operational parameters affected by the schedule, is probably not an accurate or complete description, as it does not take into account correlations between parameters, nor the varied importance of individual parameters. Nevertheless, this study shows that, while optimization of equipment operation can provide savings, it is important to ensure high reliability of implementation and to ensure that carefully thought out measures are ultimately implemented.

In practice, completing a thorough analysis of this type during the

Fig. 7. Median, upper, and lower inter-quartile and maximum and minimum total energy related costs for each probability of failure and complexity level.

Fig. 8. Impact of reliability on mean total costs.

Fig. 6. Demand vs. energy cost for each of 100 realizations of the S0-S2 optimization strategies, at various levels of communication success probability. By comparison, the cost for full nightly TES charge is shown by the red square, and the cost for optimized schedules with perfect delivery and perfect load forecast by the green triangle. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

planning of a building control would not be practical, but potentially some insight could be collected from various similar studies and guidelines developed that at a minimum avoid costly mistakes.

Acknowledgments

The work reported here has been financed in part 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 with the goal of demonstrating ultra-high-efficient buildings and microgrids with local generation and control. Additional funding was provided by the Public Service Company of New Mexico.

Appendix A

TRNSYS HVAC model

The modular structure of TRNSYS allows the assembly of a virtual system by connecting components chosen from an extensive library. Each component is described by an appropriate set of algebraic and differential equations that provide an accurate representation of the corresponding physical system. Parameters that represent the specific characteristics of a specific component (e.g., the area of a solar collector or the thermal loss coefficient of a heat storage tank) are assigned based on experimental measurements or manufacturer specifications. Because of its characteristics, TRNSYS is particularly well suited to the simulation of complex systems with highly transient nature [41]. The thermal storage tanks are represented using Type 534 modules, using 50 nodes for both the cold storage and the hot storage. Two inlet ports and two outlet ports (one of each at the top and bottom of the tank) allow charging and discharging of the thermal storage. For example, the cold storage tank is charged by a flow of chilled water entering at the inlet port corresponding to the lowest node in the tank, with a corresponding flow exiting the tank from the highest node. On the other hand, the hot storage is charged by inserting hot water at the top, and removing a corresponding amount at the bottom. Thermal losses and temperature diffusion is treated by the Type 534 module. The reader is referred to the literature for details [21]. A solar array, modeled using a combination of Type 539 flat plate collectors and Type 71 vacuum tube collectors, provides heat to a Type 5b heat exchanger module. Flow through the solar array is activated when solar irradiance is above 150 W/m² and flow rate is controlled by a Type 23 PID controller to maintain a constant array outlet temperature of 95 °C. On the storage side of the heat exchanger, flow is activated when the array outlet temperature exceeds 90 °C, and a PID controller maintains the hot storage supply temperature at a setpoint of 90 °C thereafter. The absorption chiller is modeled using a set of equations designed to mimic the performance indicated in the manufacturer's technical documentation. The flow rate of the heating medium serving the absorption chiller is constant, as is the flow rate of the chilled water and the cooling water.

The inlet temperature of the cooling water is constant at 25 °C, at the nominal flow rate of 10 kg/s. The heating medium inlet temperature is the temperature of water drawn from the hot storage, while the outlet temperature is calculated based on a 6.7 °C temperature reduction for the chilled water. The chiller is activated when called on by the schedule, under the condition that sufficient hot water is present in the hot tank (3000 kg at a temperature over 80 °C). Once activated, the chiller continues to run until called upon by the schedule, or until the hot water storage is exhausted.

The cold storage tanks are charged during off-peak hours, via a Type 5b heat exchanger. Chilled water at 6 °C (supplied by a district cooling system) at a flow rate of 13.9 kg/s enters the cold side. Water is drawn from the top outlet of the tank, cooled by the heat exchanger and returned to the bottom inlet of the tank. If the absorption chiller operates during the charging period, the flow rate to the tank inlet is increased and the temperature of the mixed stream is calculated. Tank charging stops when the tank outlet temperature drops below 10 °C, independently of the charge schedule. Fig. A1 shows the thermal profile of the chilled water tank over a period of 1 week. Temperature sensors labeled T1-T8 are providing reading of the temperature in degrees C at 8 locations from the top of the tank (T1) to the bottom (T8).

During the day, the tanks are depleted by drawing chilled water from the bottom, and serve the cooling load via a Type 508c cooling coil. The cooling coil model calculates coolant flow rate and exit temperature resulting from cooling air from the return temperature of 25 °C to a supply temperature of 15 °C, at a flow rate determined to match the cooling load. If the chiller is operating, the flow rate of coolant drawn from the tank is reduced correspondingly, and the temperature of the mixed stream from the tank and from the chiller is calculated. The cooling costs required to

Fig. A1. Example cold water tank temperature profile.

meet the cooling load that vary as a function of operating schedule are calculated based on the electricity tariff described earlier. Energy charges result from operation of the electric chiller to either charge the tanks or meet cooling load not met by either storage or absorption chiller, and from operation of various pumps associated with operating the absorption chiller (cooling tower, chilled water and heating water pumps). In addition, demand charges from operating the equipment are monitored. The maximum demand charge for each billing period is applied to the total energy bill.

References

- [1] Jones CB, Robinson M, Barsun H, Ghanbari L, Mashayekh S, Feng W, et al. 'Software-as-a-service optimal scheduling of New Mexico Buildings. ECEEE 2015 summer study on energy efficiency. Toulon/Hyères (France): Club Belambra Les Criques, Presqu'île de Giens; 2015.
- [2] Koeppl G, Andersson G. Reliability modeling of multi-carrier energy systems. *Energy* 2009;34(3):235–44.
- [3] Marnay C, Venkataramanan C, Stadler M, Siddiqui A, Firestone R, Chandran B. Optimal technology selection and operation of microgrids in commercial buildings. *IEEE Trans Power Syst* 2008;23(3). <https://doi.org/10.1109/TPWRS.2008.922654>.
- [4] GAMS. General algebraic modeling system user's guide; 2018. Available at <https://www.gams.com/latest/docs/>.
- [5] HOMER. Hybrid optimization of multiple electric renewables Pro 3.11 user manual; 2017. Available at <https://www.homerenergy.com/support/docs/3.11/index.html>.
- [6] Marnay C, Blanco R, Hamachi KS, Kawann CP, Osborn JG, Rubio FJ. Integrated assessment of dispersed energy resources deployment. Lawrence Berkeley National Laboratory; 2000. [report LBNL-46082].
- [7] Marnay C, Rubio FJ, Siddiqui AS. The shape of the microgrid. Summary of talk presented in the role of distributed generation in reinforcing the critical electric power infrastructure panel at the 2001 IEEE power engineering winter meeting, Columbus, OH, 31 January. *IEEE Power Engineering Review, Proceedings (Cat. No.01CH37194)*, 1. 2001. p. 150–3.
- [8] Siddiqui AS, Marnay C, Hamachi KS, Rubio FJ. 'Customer adoption of small-scale on-site power generation. Paper presented at the European council for an energy efficient economy summer study 2001, Mandelieu, France, 11–15 June. 2001.
- [9] Siddiqui AS, Firestone RM, Ghosh S, Stadler M, Edwards JL, Marnay C. Distributed energy resources customer adoption modeling with combined heat and power applications. Lawrence Berkeley National Laboratory; 2003. [report LBNL-52718].
- [10] Mancarella P. Distributed multi-generation and district energy systems. In: Keirstead J, Shah N, editors. *Urban energy systems: an integrated approach*. London: Routledge; 2013. p. 76–95.
- [11] Liu P, Feng W, Marnay C, Dutton S, Jin M, Zheng L, Zhou N. Towards the optimal development of low-carbon community energy systems. Paper presented at the 2016 ACEE summer study on energy efficiency in buildings, Asilomar Conference Center, Pacific Grove, CA, 21–27 Aug 2016. 2016.
- [12] Jin M, Feng W, Liu P, Marnay C, Spanos C. MOD-DR: microgrid optimal dispatch with demand response. *Appl Energy* 2017;187:758–766<https://www.sciencedirect.com/science/article/pii/S030626191631724X>.
- [13] Jin M, Feng W, Marnay C, Spanos C. Microgrid to enable optimal distributed energy retail and end-user demand response. *Appl Energy* 2018;210:1321–55<https://www.sciencedirect.com/science/article/pii/S0306261917306062>.
- [14] Stadler M, Aki H, Firestone R, Lai J, Marnay C, Siddiqui A. Distributed energy resources on-site optimization for commercial buildings with electric and thermal storage technologies. Paper presented at the 2008 ACEE summer study on energy efficiency in buildings, scaling up: building tomorrow's solutions, Asilomar, CA, August 17–22. 2008.
- [15] Stadler M, Kloess M, Groissböck M, Cardoso G, Sharma R, Bozchalui MC, et al. Electric storage in California's commercial buildings. *Appl Energy* 2013;104:711–22https://ac.els-cdn.com/S0306261912008264/1-s2.0-S0306261912008264-main.pdf?_tid=16efca87-8014-4c93-90cf-b918da5e252e&acdnat=1531873222_bc4dfe476cf7e9025a6a0551307a25f.
- [16] DeForest N, Mendes G, Stadler M, Feng W, Lai J, Marnay C. Optimal Deployment of thermal energy storage under diverse economic and climate conditions. *Appl Energy* 2014;119:488–96https://ac.els-cdn.com/S030626191400066X/1-s2.0-S030626191400066X-main.pdf?_tid=caa79b62-bfaa-4e45-9a2e-445ae6cfbc6&acdnat=1531872426_619d07feac690ab23e6adaf00b87f05.
- [17] Steen D, Stadler M, Cardoso G, Groissböck M, DeForest N, Marnay C. Modeling of thermal storage systems in MILP distributed energy resource models. *Appl Energy* 2015;137:782–92https://ac.els-cdn.com/S0306261914007181/1-s2.0-S0306261914007181-main.pdf?_tid=852eff37-df18-46c8-a895-5d8190c9caea&acdnat=1531872101_de6911b9b31628b94f8cf497d2020a.
- [18] Ortiz M, Barsun H, He P, Vorobieff P, Mammoli A. Modeling of a solar-assisted HVAC system with thermal storage. *Energy Build* 2010;42(4):500–9.
- [19] Mammoli A, Vorobieff P, Barsun H, Burnett R, Fisher D. Energetic, economic and environmental performance of a solar-thermal-assisted HVAC system. *Energy Build* 2010;42(9):1524–35.
- [20] Mammoli A, Stadler M, DeForest N, Barsun H, Burnett R, Marnay C. 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. 2013.
- [21] Kleinbach EM, Beckman WA, Klein SA. Performance study of one-dimensional models for stratified thermal storage tanks. *Sol Energy* 1993;50(2):155–66.
- [22] Cardoso G, Stadler M, Siddiqui AS, Marnay C, DeForest N, Barbosa-Póvoa A, et al. Microgrid reliability modeling and battery scheduling using stochastic linear programming. *J Electr Power Syst Res* 2013;103:61–9.
- [23] Marnay C, Aki H, Hirose K, Kwasinski A, Ogura S, Shinji T. Japan's pivot to resilience: how two microgrids fared after the 2011 earthquake. *IEEE Power Energy Mag* 2015;15(3):44–57.
- [24] Stadler M, Groissböck M, Siddiqui A, Cano EL, Marnay C. Optimizing distributed energy resources: innovative operational and strategic decision models. *Appl Energy* 2014;81:416–23https://ac.els-cdn.com/S0306261914007235/1-s2.0-S0306261914007235-main.pdf?_tid=e927c72b-dd51-481f-afcc-fe82c0472b20&acdnat=1531871694_f9d51c4ff32a271ff9b915bde3e6693c.
- [25] Nemat N, Braun M, Tenbohlen S. Optimization of unit commitment and economic dispatch in microgrids based on genetic algorithm and mixed integer programming. *Appl Energy* 2018;210:944–63.
- [26] Li Y, Feng B, Li G, Qi J, Zhao D, Mu Y. Optimal distributed generation planning in active distribution networks considering integration of energy storage. *Appl Energy* 2018;210:1073–81.
- [27] Quashie M, Marnay C, Bouffard F, Joós G. Optimal planning of microgrid power and operating reserve capacity. *Appl Energy* 2018;210:1229–36<https://www.sciencedirect.com/science/article/pii/S0306261917310309>.
- [28] Wang C, Yan J, Marnay C, Djilali N, Dahlquist E, Wu J, et al. Distributed Energy and Microgrids (DEM). *Appl Energy* 2018;210:685–9https://ac.els-cdn.com/S0306261917316550/1-s2.0-S0306261917316550-main.pdf?_tid=5721989b-dbdd-4284-8594-83182ebf5c80&acdnat=1531870650_b1cb8a99e47003b577fc68b5c4a6d6da.
- [29] Gao DC, Sun Y, Lu Y. A robust demand response control of commercial buildings for smart grid under load prediction uncertainty. *Energy* 2015;93:275–83.
- [30] Wang L, Mathew P, Pang X. Uncertainties in energy consumption introduced by building operations and weather for a medium-size office building. *Energy Build* 2012;53:152–8.
- [31] Ahmadi A, Nezhad AE, Hredzak B. Security-constrained unit commitment in presence of lithium-ion battery storage units using information-gap decision theory. *IEEE Trans Ind Inf* 2018.
- [32] Kazemi M, Zareipour H, Amjadi N, Rosehart WD, Ehsan M. Operation scheduling of battery storage systems in joint energy and ancillary services markets. *IEEE Trans Sustain Energy* 2017;8(4):1726–35.
- [33] Najafi-Ghalelou A, Nojavan S, Zare K. Information gap decision theory-based risk-constrained scheduling of smart home energy consumption in the presence of solar thermal storage system. *Sol Energy* 2018;163:271–87.
- [34] Najafi-Ghalelou A, Nojavan S, Zare K. Heating and power hub models for robust performance of smart building using information gap decision theory. *Int J Electr Power Energy Syst* 2018;98:23–35.
- [35] Mehdizadeh A, Taghizadeh N, Salehi J. Risk-based energy management of renewable-based microgrid using information gap decision theory in the presence of peak load management. *Appl Energy* 2018;211:617–30.
- [36] Stadler M, Marnay C, Donadee J, Lai J, Mégel O, Bhattacharya P, Siddiqui A. Distributed energy resource optimization using a Software as Service (SaaS) approach at the University of California, Davis Campus. Lawrence Berkeley National Laboratory; 2011. [report LBNL-4285E].
- [37] Manur A, Venkataramanan G, Sehloff D. Simple electric utility platform: a hardware/software solution for operating emergent microgrids. *Appl Energy* 2018;210:748–63.
- [38] Zheng M, Wang X, Meinrenken C, Ding Y. Economic and environmental benefits of coordinating dispatch among distributed energy storage. *Appl Energy* 2018;210:842–55.
- [39] DeForest N, MacDonald JS, Black DR. Day ahead optimization of an electric vehicle fleet providing ancillary services in the Los Angeles Air Force Base vehicle-to-grid demonstration. *Appl Energy* 2018;210:987–1001.
- [40] Hale PS, Arno RG. Survey of reliability and availability information for power distribution, power generation, and HVAC components for commercial, industrial, and utility installations. IEEE industrial and commercial power systems technical conference. Conference record (Cat. No.00CH37053). 2000. p. 31–54.
- [41] Beckman WA, Broman L, Fiksel A, Klein SA, Lindberg E, Schuler M, et al. TRNSYS the most complete solar energy system modeling and simulation software. *Renew Energy* 1994;5(1–4):486–8.