

DEMAND RESPONSE ASSESSMENT TOOL: A CLOUD-BASED SIMULATION TOOL FOR RAPID ASSESSMENT OF DEMAND RESPONSE POTENTIAL IN COMMERCIAL AND INSTITUTIONAL FACILITIES

Benjamin Futrell, Madison Wynn, Eric Tate, and Robert W. Cox
 Energy Production and Infrastructure Center
 University of North Carolina Charlotte
 Charlotte, NC USA

ABSTRACT

Mid-to-large size commercial buildings are not significant participants in demand-response programs because of the risk and effort associated with predicting the actual amount of load they can shed. This paper describes a tool designed to rapidly predict the amount of demand reduction these buildings can provide. The tool uses commercial building reference models published by the Department of Energy. Users can adjust certain key parameters so that peak demands given by the reference model match actual billed demand. The paper describes the basic tool and provides field test results from several schools in El Paso, Texas.

INTRODUCTION

The Federal Energy Regulatory Commission (FERC) defines demand response (DR) as "a reduction in the consumption of electric energy by customers from their expected consumption in response to an increase in the price of electric energy or to incentive payments designed to induce lower consumption." (FERC 2008). Figure 1 summarizes the impact of DR programs implemented by utilities and grid operators in the United States in 2018 (EIA 2019b; EIA 2019a). Note that the 9.75 million participating customers were given nearly \$1.2B in incentive payments and bill savings in exchange for a peak demand reduction of 12,522MW. This load savings is equivalent to the output of roughly 160 peaking power plants (GAO 2014).

Utilities and grid operators have a statutory requirement to provide reliable electric power, so voluntary demand reductions from customers are not necessarily the most natural choice for managing peak loads and grid stability. Given their need for certainty when calling for a demand reduction, utilities and grid operators prefer to work with industrial customers. This is clear in Figure 1, which shows that about half of the overall peak demand reduction came from industrial customers who represent less than 1% of the overall participation base. Industrial plants are ideal because they are large and only a handful of them are required to have a significant impact. More impor-

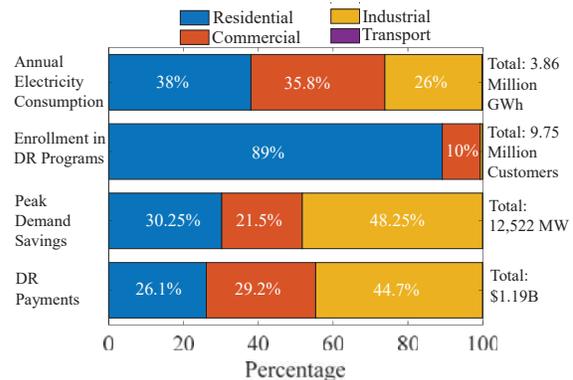


Figure 1: Total electricity consumption and DR program participation by sector in the United States in 2018.

tantly, many large industrial customers have backup generators and utilities can reliably call upon these customers to utilize these generators to help reduce the demand on the grid during peak periods for an appropriate incentive payment. The residential sector also provides a relatively reliable source of demand reduction by turning off air conditioners and water heaters during peak periods. Utilities minimize the risk of underestimating the amount of demand response by enrolling a large number of small independent customers, so that the relative impact of any one failure is negligible. For this reason, nearly 90% of DR participants come from the residential sector.

Figure 1 indicates that commercial buildings rank third in their contribution to peak demand savings despite consuming nearly the same amount of electricity as the residential sector. Understanding why the commercial sector lags behind requires consideration of the unique characteristics of this building set. Most importantly, the commercial sector is extremely diverse, ranging from small retail outlets that differ only slightly from single-family homes to large hospital campuses that rival industrial plants. Large commercial customers occasionally provide backup generation, and research has demonstrated

significant demand-reduction potential in small commercial buildings having only single-stage rooftop units (Cai and Braun 2019).

Mid-to-large size commercial buildings also present significant DR opportunities, but their complexity makes it more difficult to predict savings. Building owners and DR aggregators both see risk in the commercial space, since building owners do not want to discomfort their revenue-generating tenants and DR aggregators do not want to have buildings underperform on demand-limiting days. For full confidence, a facility audit and calibrated energy model must be created. Some predictive tools have been developed and data-driven approaches have been presented, but these have not been widely adopted. This paper builds upon an approach first presented in (Yin et al. 2010), which used prototypical reference models and a limited set of tunable input parameters to predict demand reductions from cooling setbacks on a limited set of peak demand days in commercial offices in California. The efforts presented here focus heavily on a calibration approach intended to avoid submetering to allow rapid assessment. The tool is also open-source and web-based, thus allowing other users the opportunity to further develop it to suit their needs. The software and approach have been tested in several offices and schools participating in DR programs in El Paso, Texas.

The remainder of this paper describes the approach developed and tested by the authors. The next section presents related work and motivates the need. The subsequent sections describe the details of the approach and present validated field results.

BACKGROUND

Researchers have long studied the potential for demand savings in commercial buildings. Significant opportunities exist because of the large and distributed thermal mass inherent in these facilities. Wang et al. (2014) provided a thorough review of potential DR strategies. The primary opportunity is to shift thermal load by precooling the building during off-peak periods and then adjusting the thermostat set-points upward during on-peak periods. Many papers have utilized simulation studies to develop optimal setpoint adjustment strategies focused either on maximizing demand savings or minimizing overall utility cost (Wang, Xue, and Yan 2014; Braun 1990; Lee and Braun 2016; Korkas et al. 2016; Xue et al. 2015). The primary goal is to ensure that savings can be achieved without sacrificing occupant comfort. Many of these studies focus on developing strategies that provide overall energy and cost savings for the customer by exploiting time-of-use rate structures.

This paper focuses on a specific need that has emerged in many electricity markets in recent years. Facilities can now generate revenue by working with DR aggregators

who can bid demand reductions into certain markets. Many energy-services companies and controls vendors are interested to participate as DR aggregators, but they want to accurately predict the amount of demand savings that can be bid. Savings can be reasonably predicted by developing calibrated energy models but creating them can be costly and time-consuming.

DR aggregators are interested in having a tool that can rapidly predict demand savings on a handful of peak demand days. Several software packages and approaches have been developed for this purpose, but all of them have certain limitations. One powerful package is the Demand-Limiting Assessment Tool (DLAT) (Purdue 2010). This software was derived from earlier simulation programs described by (Mercer and Braun 2005) that were developed for modeling the dynamic energy performance of small commercial buildings. This package performs relatively detailed analysis, including transfer function models for building walls and was validated by comparisons with Energy-10 and TRNSYS (Mercer 2003). In its current form, DLAT is designed for broad studies of small commercial buildings. It only incorporates single-stage rooftop units with compressors that cycle on and off to meet the building cooling load. Furthermore, it provides no ability to adjust the building envelope and orientation. According to the accompanying help files, DLAT can represent minimally code-compliant buildings for several different vintages of Title 24 as well as ASHRAE 90.1-2007 and a typical 1985 construction (Purdue 2010). DLAT can represent seven types of small buildings, but each of these do not have much flexibility in terms of key parameters. For instance, the ASHRAE 90.1-2007 models each have fixed window-to-wall ratios that depend on the building type and each has windows with a solar heat gain coefficient (SHGC) of 0.49. In many southern parts of the United States, for instance, the prescribed SHGC values are much lower (ASHRAE 2007). DLAT was found to be too constrained for use in our field studies.

The Demand Response Quick Assessment Tool (DRQAT), developed by Lawrence Berkeley National Laboratory, is another software package that can predict load-shedding potential (LBNL 2010). This free software provides more flexibility than DLAT, allowing the user to alter key aspects of the building envelope and geometry and to introduce either water-cooled or air-cooled chillers. DRQAT uses the DOE's Large, Medium, and Small Office Reference Models and simulates these using Energy Plus. The user can modify key inputs such as the window-to-wall ratio, building orientation, and loads such as lighting and equipment power density. (Yin et al. 2010) demonstrated that DRQAT was able to reasonably predict savings on DR event days in a series of experimental tests performed in California

offices. The approach presented by (Yin et al. 2010) is appealing, but it is somewhat limited because only office buildings in California can be simulated, and it cannot support significant space-type diversity in a given building. Additionally, it does not allow one to modify the efficiencies of the cooling equipment, which can vary significantly based on building type.

This paper builds upon the successful approach demonstrated by DRQAT. The software developed here focuses on schools, commercial offices, and retail sites with an emphasis on climate zones and envelopes found outside of California. This paper includes field studies validating the performance of the software in schools with significant space-type diversity and loading patterns that vary significantly because of summertime occupancy. The diversity issues seen here are far greater than those described in (Yin et al. 2010). Those researchers also demonstrated the effectiveness of a calibrated reference model for predicting demand savings in offices, and their approach required submetering. In this case, we were interested in developing a calibration procedure that did not require any submeters. The software described here is also open-source and uses a cloud-based platform that can be easily updated by any qualified user. Contributors from different regions can easily add building types and climate zones as needed.

SYSTEM DESCRIPTION

The concept behind the proposed Demand Response Assessment Tool is that a generic set of reference models can be appropriately used to quickly determine the potential amount of load that can be shed by adjusting thermostat setpoints in a mid-to-large size commercial building.

Figure 2 shows the basic arrangement. A Linux-based compute server hosted on a cloud platform runs several scripts that can take user inputs and use them to start an energy simulation in EnergyPlus. Users interface with this system through a web-based dashboard. Basic building inputs are provided via the web interface and pushed to the server. When the user is ready to perform analysis, an event-detection script running on the server calls EnergyPlus and generates a new set of energy and demand estimates. Final results are emailed to the dashboard, which summarily presents them to the user. Results are also saved to a file repository located on the compute server. This architecture allows the system to be easily tested and updated without having to push new versions to operators in the field or having to invest in the development of a standalone application. This decision was important since the authors wanted to create a "living" tool that could be easily edited by multiple contributors without requiring extensive development cycles.

Figure 3 shows the user interface, which is displayed in a web browser. The left side includes entries for input parameters, and the right side shows results obtained from

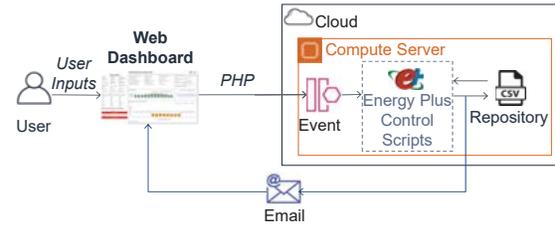


Figure 2: Architecture of the proposed Demand Response Assessment Tool (DRAT).

the energy simulations.

Figure 4 shows the basic process for using the Demand Response Assessment Tool. First, the user provides basic information about the candidate building, such as its location and footprint size. Second, the user must provide demand information obtained from electricity bills over the last twelve months. Detailed interval meter data is not required. Finally, users begin a process of "demand calibration" that focuses on modifying the model to match the billed demand over the last twelve months. The following sections describe each of these basic steps.

Step 1: Enter Basic Building Information

Figure 5 shows the parameters the user must provide. The building type dropdown menu selects one of the Department of Energy's commercial building reference models (Michael Deru and Torcellini 2011). Although there are 16 such models available, the authors have only field-tested the medium office and secondary school. The basic operating principle is to use these reference models and to adjust them according to the information available about the target building.

The software uses the basic geometry of the DOE reference models and allows the user to modify certain critical inputs. Figure 5 shows that the user first selects the climate zone and city from a dropdown menu. This choice controls the weather file selected when the model is simulated. The user then provides key information about the building's envelope, HVAC system, and individual load components. The material properties of the building envelope are modified by selecting the appropriate energy code and city. Currently, the system includes material properties specified in the 1989, 2004, and 2007 versions of ASHRAE 90.1. The system also includes a set of parameters identified in the "pre 1980" models created by DOE. Specific parameter values are also updated when the user changes the climate zone. For instance, the solar-heat gain coefficient of the glass is updated appropriately when the user changes either the energy code or the climate zone. Note that one can also specify the window-to-wall ratio, which can vary significantly between buildings.

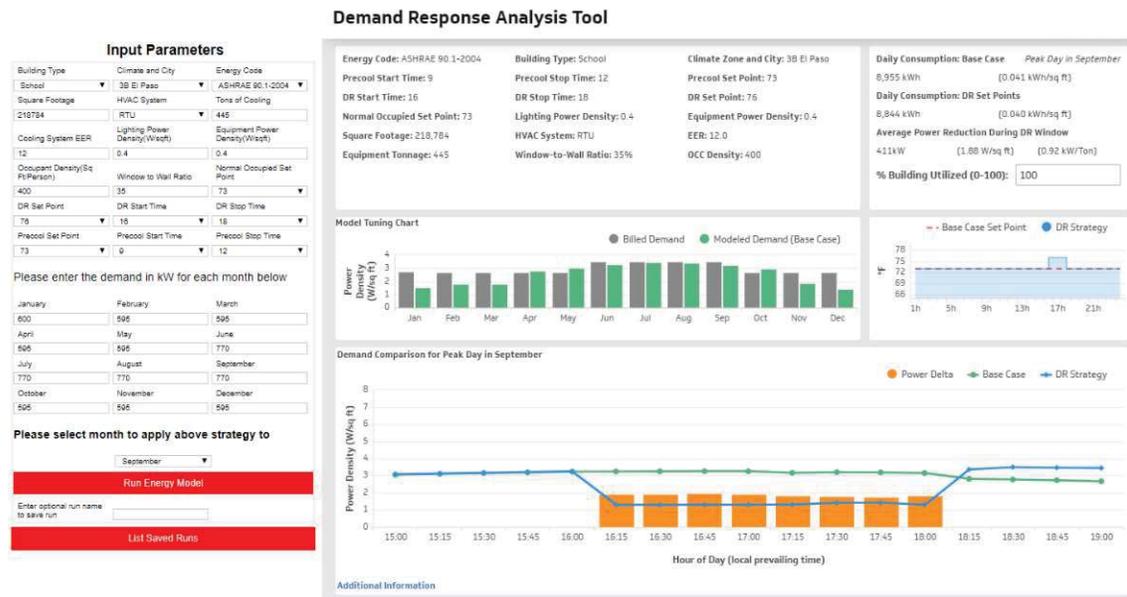


Figure 3: The current user interface. The left side includes input parameters about the building, and the right side features results obtained from the energy simulations.

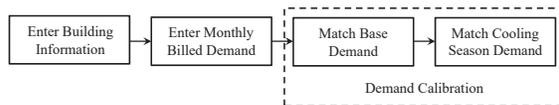


Figure 4: Basic steps for utilizing DRAT to rapidly match the reference model to billed demand.

Figure 5 shows that the user can also specify information about the HVAC system. All building types are assumed to use variable-air volume (VAV) systems, since this is typical for mid-to-large size commercial buildings. The user can select the primary source of conditioning using the "HVAC System" menu. The choices include rooftop units (RTUs), a water-cooled chiller, and an air-cooled chiller. Users must also provide the total building cooling capacity in tons and the total conditioned floor area in square feet. To keep the assessment process simple, the software does not modify the size of the reference models. Instead, it takes the user input and computes the cooling capacity per unit area. The cooling capacity in the reference model is then updated so that the cooling capacity per unit area in the model matches that of the real building.

Figure 5 also shows that the user can specify the building equipment and occupancy loads. Entry fields are provided for the equipment power density, lighting power density, and occupant density. A user manual provides basic guidance on selecting these values based on the energy code

and typical diversity factors. For instance, the lighting power and equipment power densities in a school differ based on the space type (i.e. classroom, hallway, gymnasium, etc.). The user must choose values taking such diversity into account. Users can also make decisions based on observations from the target building itself. The software uses the schedules included in the reference models, but the schedules tend not to be as important as in a typical energy model since the focus is to match the peak demand.

The last set of simulation parameters to be entered in the dashboard are those that control the space-temperature setpoints before, during, and after a DR event. Figure 5 shows that a user can enter the following key values, including the normal occupied cooling setpoint, the setpoint during the DR event, the start time for the DR event, the end time for the DR event, and the setpoint and timing for an optional pre-cool period.

Step 2: Enter Monthly Billed Demand and Execute Initial Simulation Run

The ultimate goal is to modify a reference model to match the peak demand in a real building. To accomplish this, the user must provide the billed demand over the last 12 months. Figure 5 shows that individual entry fields are provided for each month.

Following any given simulation run, the results presented on the right side of Figure 3 are updated. Key areas include the "Model Tuning Chart" displayed in the middle

Input Parameters

Building Type	Climate and City	Energy Code
School	3B El Paso	ASHRAE 90.1-2004
Square Footage	HVAC System	Tons of Cooling
218784	RTU	445
Cooling System EER	Lighting Power Density(W/sqft)	Equipment Power Density(W/sqft)
12	0.4	0.4
Occupant Density(Sq Ft/Person)	Window to Wall Ratio	Normal Occupied Set Point
400	35	73
DR Set Point	DR Start Time	DR Stop Time
76	18	18
Precool Set Point	Precool Start Time	Precool Stop Time
73	9	12

Please enter the demand in kW for each month below

January	February	March
600	595	595
April	May	June
595	595	770
July	August	September
770	770	770
October	November	December
595	595	595

Figure 5: Details of the parameters used to describe the target building.

of the screen. The gray bars show the billed demand for each month as provided by the operator; the green bars show the modeled demand. The bottom chart, which is titled the "Demand Comparison for the Peak Day", compares the expected whole-building power density during the DR event to what it would have been had the setpoints remained at their normally occupied values. The chart also shows the difference between these two time series in each 15-minute window during the DR event.

Step 3: Demand Calibration

The last step is to tune the model so that billed peak demand and modeled peak demand match during the cooling season when DR events are likely to occur. This tuning uses actual measured weather data recorded during the calibration window. Figure 4 shows that this tuning occurs in two basic steps. First, the user adjusts the baseload parameters, namely the equipment power density and lighting power density. These values are adjusted to account for the diversity of space types throughout the target building. Specific choices can be chosen using a combination of energy codes and field observations. During the calibration phase, the user adjusts the values so that modeled and billed demand match during cooler months when cooling loads are minimal. The goal is to match the peak afternoon demand, even in the winter. In the climate zones (3A and 3B) and building types (medium office and school) tested thus far, there are months when the afternoon peak is entirely satisfied by economizing. The user

can thus be confident that equipment and lighting power density are the primary drivers of peak demand during these months and they can be adjusted accordingly. For a building with electric heating, submeters might be needed.

Once the baseload is tuned, the final step is to adjust the efficiency of the HVAC system so that the peak demand matches during the cooling season. Specifically, the user adjusts the energy efficiency ratio (EER) of the HVAC equipment so that the whole-building peak demand matches closely during the cooling season. Note that such an adjustment is reasonable for several reasons. First, it is likely that the installed EER differs from the nameplate value both as a result of age and system operational issues. Second, the HVAC system included in the model must represent the entire set of space conditioning for the building. Many commercial buildings, however, include a primary system and a set of smaller systems. Schools, for instance, may have small areas served by rooftop units even though the vast majority of the space is served by an air-cooled chiller. This aggregate EER allows one to reflect this small diversity of systems. The user can also tune the window-to-wall ratio to help adjust the cooling load. In general, this value can be determined using photographs publicly available on the internet.

Please note that the authors could focus on performing calibration entirely consistent with the metrics in ASHRAE Guideline 14 (ASHRAE 2014). Since the goal of this tool is only to predict the demand on the peak cooling days, the authors were not focused on satisfying Guideline 14 targets for energy. In fact, the authors could not satisfy those requirements without requiring more detailed inputs about operating schedules. This issue is particularly problematic in schools, for which different schedules are needed for different parts of the year. Had only offices been considered, it may have been possible to more fully calibrate the model to match energy consumption.

The authors note the heuristic nature of the proposed approach, but provide several caveats. First, the predicted peak demand values have all been within $\pm 10\%$ of their true values in all valid experimental cases thus far. As the next section shows, the authors have also obtained low values for both the MBE and CVRMSE for demand based on monthly data. ASHRAE Guideline 14 does provide guidance on demand calibration when using whole-building energy models, but the values obtained here are inline with what would be expected if the model met the energy acceptance criteria. Furthermore, the authors could calibrate using hourly interval data, but this was not possible in the 2019 cooling season. This approach will be further investigated and compared to the current approach in the 2020 cooling season.

FIELD TESTS

Our approach was tested at several buildings in El Paso, Texas. This section describes two key sets of initial tests. First, it describes a set of updates that were required for the DOE reference models. Second, it provides an illustrative case study from a high school in El Paso.

Initial Model Tuning

One well known challenge associated with predicting the effectiveness of setback-based DR programs is the ability to properly model thermal mass. To investigate the significance of this issue, initial tests were performed using an office building in Durham, North Carolina. Figure 6 shows the results from this study. The top graph shows the measured results on a day when the high temperature was 36.1°C. The yellow, blue, and green curves show the behaviors in the hottest, average, and coolest zones respectively. Space conditioning was provided until 11AM, when the zone setpoints were setback to 29.4°C. The marker shows that the maximum rate of change in the hottest zone was about 1.1°C/hr.

The middle graph shows the initial behavior of the commercial building reference model provided by DOE, again showing the behaviors of the hottest, average, and coolest zones. Cooling was also terminated at 11AM. Note that the zone temperatures initially increased at rates as high as 10.1°C/hr, and temperatures consistently rose faster than the maximum value observed in the real building (1.1°C/hr).

The bottom graph shows the behavior of the commercial building reference model after increasing the thermal mass. Note that the initial rate of change is still high over the first 15 minutes, but eventually becomes more gradual. The initial rise appears to be an artifact of the equation solvers in EnergyPlus. Note that both the low-mass and high-mass models had essentially the same annual energy consumption and both predicted the peak demand within $\pm 10\%$. When attempting to predict the demand savings, however, the low-mass model was off by more than 50% since its space conditioning systems turned on almost immediately after the start of a DR event.

Figure 6 highlights a key difficulty associated with using simulation models, particularly reference models, to predict the load-shedding capability in a commercial building. (Yin et al. 2010) identified similar issues with DRQAT. The results presented in the next section suggest that the heuristic tuning method utilized in this paper can still provide reasonably accurate predictions. One could potentially consider approaches such as those described in (Lee and Hong. 2017) to tune the uncertain thermal mass using zone-temperature measurements from a building automation system.

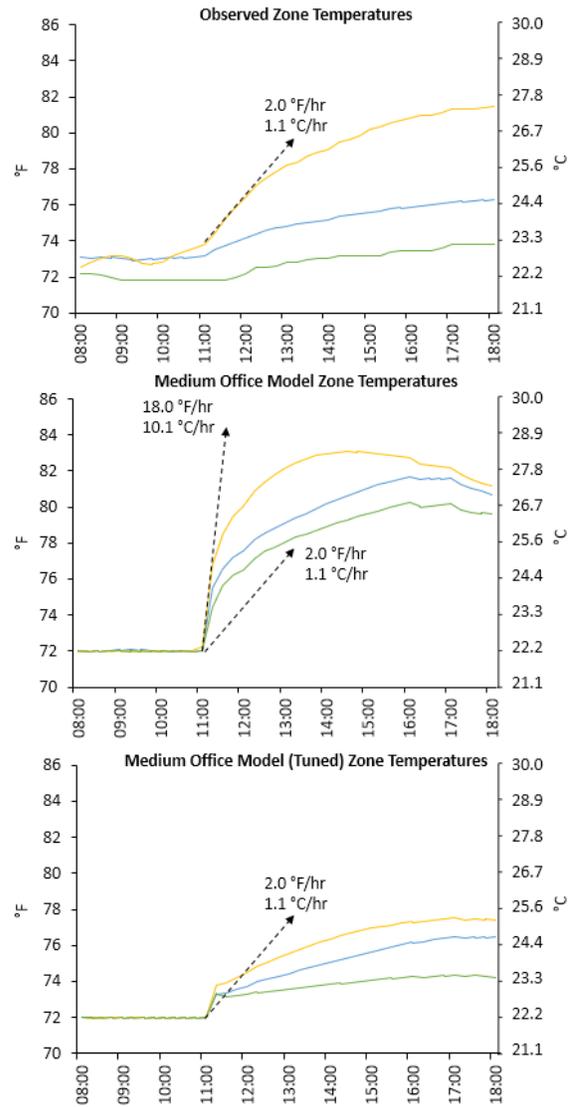


Figure 6: Modeled and observed temperature patterns when space-conditioning systems are disabled on hot days. Top: Field measurements. Middle: Initial results from the DOE medium office model during the same weather conditions. Bottom: Model results after tuning the thermal mass parameters. The yellow, blue, and green curves show the temperatures in the hottest, average, and coolest zones, respectively.

Example Results in a High School in Texas

The authors worked with their collaborators to test the approach in several sites near El Paso, Texas. Testing was performed in buildings enrolled in the Transmission/Distribution Service Provider (TDSP) Load Manage-

ment Program administered by the local utility in response to instructions issued by the Electric Reliability Council of Texas (ERCOT).

The example building is a 20,326m² high school served by rooftop units having an aggregate capacity of 1.559MW (445 tons). Given the age of the building, it was assumed to be built roughly to the standards included in ASHRAE 90.1-2004. The occupant density was set to 18.6ft²/person since the school has approximately 1000 total students.

Figure 7 shows how the demand data is used to tune the model. The top graph shows the results after the initial model run, the middle graph shows the results after tuning the base loads, and the bottom graph shows the results after tuning the HVAC efficiency and window-to-wall ratio. Red boxes are included to highlight the values targeted as the process progresses from one chart to the next. Initially, the user focuses on tuning the baseload, and the emphasis is placed on January, February, November, and December. During these months, the building is primarily economizing so the peak load is driven by equipment and lighting. After adjusting these values, the modeled and billed demand match relatively closely, with some underprediction in December. To obtain the results shown in the bottom graph, the user focuses on adjusting the HVAC system efficiency and considers the cooling months. In this case, emphasis is placed on April, May, August, and September. June and July are neglected because school is not in session. If one desires to predict the performance in these months, additional runs could be constructed, but the occupant density, lighting, and equipment power would all need to be adjusted accordingly. Note the close agreement in the cooling months in the bottom graph, particularly in August and September when the utility is likely to call for demand reductions. Note also that the final MBE and CVRMSE values for the demand are 5.2% and 15.1%, respectively. As noted previously, Guideline 14 does not provide acceptance criteria for demand when using calibrated energy models. Note, however, that these values for demand error are lower than most models would achieve if meeting the Guideline 14 acceptance criteria for energy.

The utility called for several DR events in August and September 2019. Figure 8 compares the measured and predicted results obtained during an event on September 18, 2019. During this event, the occupied setpoints throughout the school were increased from 22.78°C to 24.44°C between 3:30PM and 5:30PM. The top graph shows interval data recorded over the entire day, and the DR window is highlighted. In its present state, our software only allows DR events to begin on the hour so a comparison simulation was run from 4:00PM to 6:00PM. The top graph shows that the average load during the event

window was about 19.91W/m². The bottom graph, on the other hand, shows that our software determines the load over the same window to be about 18.94W/m², representing a prediction error of about 5%. Note that DLAT predicts a peak load of only about 10.76W/m², which is about 46% lower than the actual.

Table 1 summarizes results from 4 high schools throughout El Paso. Note that the results are within ±10% in the first 3 cases, but significantly off in the last case. The error found at school number 4 highlights the risk associated with DR in commercial buildings. In that case, the maintenance staff had overridden automatic controls on the air-cooled chiller and the system thus did not respond appropriately when called to shed load.

CONCLUSION

This paper has described a tool designed to quickly assess the amount of demand-reduction potential in a mid-

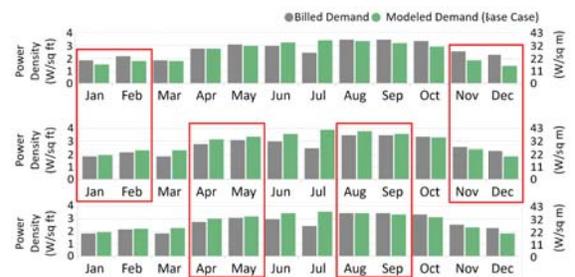


Figure 7: Progression of the modeled and billed demand as the model is tuned. Top: Results following initial run. Middle: Results following the baseload tuning step. Bottom: Results following the HVAC efficiency tuning step. The red rectangles highlight the areas that change as one moves from top to bottom.

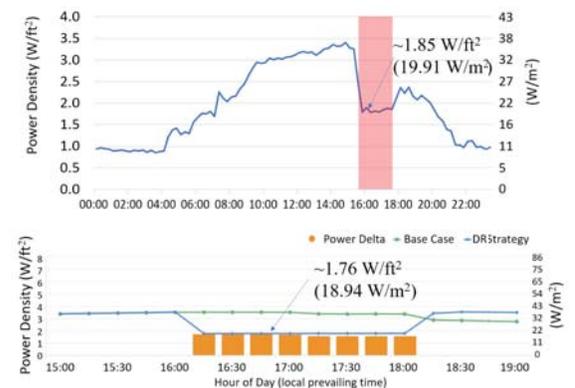


Figure 8: Measured and modeled results in the example high school during a DR event on September 18, 2019. Top: Measured results. Bottom: Modeled results.

to-large commercial building. The basic methodology is to modify the loads in a DOE reference model to match the loads in the target building. The approach is demonstrated in high schools in El Paso, Texas.

Table 1: Results from several schools in the El Paso area on September 18, 2019.

School Number	Predicted Load During DR Period (W/m ²)	Actual Load During DR Period (W/m ²)
1	18.94	19.91
2	12.06	11.63
3	10.01	9.15
4	11.84	5.38

The results presented in this paper demonstrate that the methodology used to develop our software can be used to reasonably predict DR potential in mid-to-large-size commercial buildings. The calibration approach is heuristic, but appears effective in these limited test cases. The approach will be further tested using interval meter data in the 2020 cooling season. Future work will investigate the foundational merit of the heuristic approach using detailed submetering.

ACKNOWLEDGMENT

This material is based upon work supported by the National Science Foundation under grant no. IIP-1161031.

REFERENCES

- ASHRAE. 2007. "ASHRAE/IESNA 90.1 Standard-2007-Energy Standard for Buildings Except Low-Rise Residential Buildings." Technical Report.
- ASHRAE. 2014. "ASHRAE Guideline 14: Measurement of Energy, Demand, and Water Savings." Technical Report.
- Braun, J.E. 1990. "Reducing energy costs and peak electrical demand through optimal control of building thermal storage." *ASHRAE Transactions* 96:876–888.
- Cai, J., and J.E. Braun. 2019. "Assessments of demand response potential in small commercial buildings across the United States." *Science and Technology for the Built Environment* 42:1437–1455.
- EIA. 2019a. Annual Electric Power Industry Report, Form EIA-861 detailed data files. <https://www.eia.gov/electricity/data/eia861/>.
- EIA. 2019b. EIA Electric Power Monthly. https://www.eia.gov/electricity/monthly/epm_table_grapher.php?t=epmt_5_01.
- FERC. 2008. FERC Order 719. <https://www.ferc.gov/whats-new/comm-meet/2008/101608/E-1.pdf>.
- GAO. 2014. Electricity Markets: Demand-Response Activities Have Increased, but FERC Could Improve Data Collection and Reporting Efforts.
- Korkas, C. D., S. Baldi, I. Michailidis, and E.B. Kosmatopoulos. 2016. "Occupancy-based demand response and thermal comfort optimization in microgrids with renewable energy sources and energy storage." *Applied Energy* 163:93–104.
- LBL. 2010. Demand Response Quick Assessment Tool (DRQAT). <https://drcc.lbl.gov/tools/demand-response-quick-assessment-tool-drqat>.
- Lee, K.-H., and J.E. Braun. 2016. "A data-driven method for determining zone temperature trajectories that minimize peak electrical demand." *ASHRAE Transactions* 114:65–75.
- Lee, S. H., and T. Hong. 2017. "Leveraging zone air temperature data to improve physics-based energy simulation of existing buildings." *Proceedings of Building Simulation 2017*.
- Mercer, K. 2003. "Modeling and testing strategies for evaluating ventilation load reduction technologies." Technical Report.
- Mercer, K., and J.E. Braun. 2005. "Evaluation of demand-controlled ventilation and enthalpy exchangers in small commercial buildings." *ASHRAE Transactions* 111:873—889.
- Michael Deru, Kristin Field, Daniel Studer Kyle Benne Brent Griffith, and Paul Torcellini. 2011. "U.S. Department of Energy Commercial Reference Building Models of the National Building Stock." Technical Report.
- Purdue. 2010. Demand-Limiting Assessment Tool. <https://engineering.purdue.edu/DLAT/documentation.html>.
- Wang, S., X. Xue, and C. Yan. 2014. "Building power demand response methods toward smart grid." *HVACR Research* 20:665–686.
- Xue, X., S. Wang, C. Yan, and B. Cui. 2015. "A fast chiller power demand response control strategy for buildings connected to smart grid." *Applied Energy* 137:77–87.
- Yin, R., P. Xu, M.A. Piette, and S. Kilicote. 2010. "Study on Auto-DR and pre-cooling of commercial buildings with thermal mass in California." *Energy and Buildings* 42:967–975.