A Field Evaluation of Cloud-Based Predictive HVAC Control Software – The Next Step in Energy Efficiency, or Marketing Mumbo Jumbo?

David J. Gilliland, kW Engineering

ABSTRACT

Direct digital control (DDC) energy management and controls systems (EMCS) have become the standard throughout California, particularly in large office buildings. A 2012 Pacific Gas and Electric (PG&E) study found that 77% of large office buildings in their service territory utilize EMCS, a percentage likely to increase since DDC EMCS is required by California's Title 24 building code. DDC systems are no longer the cutting edge of energy efficiency – they are the new standard. So what comes next? What technology will bring efficiency beyond the capabilities of the current EMCS?

One potential avenue for increased energy savings are software-as-a-service (SaaS) offerings that remotely optimize building energy use beyond traditional EMCS capabilities. These cloud-based software solutions use statistical models to predict building loads and fine tune HVAC equipment set points to arrive at their most efficient operating parameters. Several of the software providers claim that their predictive controls yield 10-40% reductions in HVAC energy cost compared to traditional EMCS.

To test these claims, we conducted an independent, multi-year measurement and verification study of one software provider's technology. Through measuring baseline and post-retrofit HVAC energy consumption, outdoor air temperature, and occupant comfort conditions, we confirmed the potential for energy savings and investigated the impact of the predictive controls on building occupants. Our study also identified challenges and potential barriers to the technology's implementation. The resulting case study provides useful information for utilities considering adopting these technologies into incentive programs, for customers considering implementing this technology, and for software providers to improve their offerings.

Introduction

California's energy efficiency goals are among the most aggressive in the nation, driving customers toward new, more energy-efficient products and systems. Over the past 10 years this push toward efficiency has caused technologies such as variable speed drives, efficient fluorescent lighting, high efficiency compressors, and direct digital control (DDC) energy management systems (EMCS) to grow significantly in market penetration. Now, many of these once cutting-edge technologies have become commonplace. Indeed, California's latest energy code, Title 24-2013, requires many of these technologies in new construction and renovation – they are the new standard (Title 24, 2013). But California's energy efficiency goals are increasingly aggressive. Senate Bill 350, passed in November 2015, seeks to double the state's previous energy efficiency goals by the year 2030 (SB 350, 2015). In order to meet these new goals, we must look to innovative new technologies that offer savings above and beyond the technologies of the past decade.

One such series of new technologies are cloud-based predictive HVAC optimization software packages. Vendors of these software packages typically offer a Software-as-a-Service (SaaS) model, with the end-user paying a monthly fee for the ongoing operation of the software. The software is implemented on top of an existing DDC EMCS and takes control of HVAC operational set points. While a traditional EMCS tries to optimize system operations based on real-time operating conditions, these predictive software packages are designed to optimize set points based not only on real-time data, but also on the predicted thermal conditions in the building throughout the day. The vendors of these cloud-based controls packages claim they can reduce a customer's annual HVAC energy consumption by 10% to 40% by adjusting set points pre-emptively based on the predicted upcoming operation of the building.

Each software vendor claims to have a unique algorithm that achieves optimal savings. Unfortunately, since the control strategies are the primary differentiators between these companies' offerings, the algorithms are often proprietary. As a result, many in the energy efficiency industry regard these predictive software packages as the quintessential 'black box' – a technology that promises results but offers very little insight in to how it works. Customers may be hesitant to adopt these products because it is not obvious how they reduce energy use, it is difficult to differentiate between alternate product offerings, and there are few independent third party evaluations that validate the claimed savings. Similarly, in the early 2000's, companies selling black box 'power conditioning' devices promised customers energy savings via improved power quality, while providing very little insight or data on how energy was saved. While some of these power conditioning devices did work, enough customers were duped by savvy marketing that Pacific Gas and Electric (PG&E), Northern California's primary Investor-Owned Utility (IOU), published a warning for its customers encouraging them to highly scrutinize these technologies before making any financial commitments (PG&E, 2004).

This is not to say that all 'black box' technologies are bad. The creators of these predictive HVAC controls have every right to keep their algorithms private, and their savings claims may very well be legitimate. If this is the case, then these technologies represent a significant opportunity to further reduce building HVAC energy consumption in buildings that already operate with a DDC EMCS. The significance of this opportunity is supported by both the prevalence of EMCS technologies in existing buildings, and the large percentage of building energy consumption associated with HVAC equipment. Based on a 2012 study conducted by PG&E, 69% of large commercial buildings in their territory use an energy management system. The same study identifies an even greater presence in large office buildings, with a 77% penetration (Pande, 2012). In California, HVAC accounts for approximately 40% of electrical energy in large office buildings (Itron, 2006). Therefore, even an incremental reduction in HVAC usage due to improved EMCS control, if implemented broadly, could yield a significant reduction in overall state-wide building energy consumption.

Due to the industry's limited understanding of how cloud-based predictive HVAC controls work, objective assessments of these technologies will be crucial to improving consumer confidence and achieving the significant state-wide savings potential that these technologies have. To test the claims of one such technology, kW Engineering conducted a multi-year measurement and verification (M&V) case study through San Diego Gas and Electric's (SDG&E) Emerging Technologies (ET) program. This study represented an incremental but important step toward determining whether or not cloud-based predictive HVAC controls could yield reliable savings above and beyond today's energy efficiency standards.

Technology Description

The following technology description represents the vendor's own claims about the energy-saving capabilities of the product that kW Engineering and SDG&E reviewed. The technology is a cloud-based SaaS platform that offers customers energy savings and demand response capabilities beyond what their existing EMCS can provide. This SaaS platform communicates directly with the existing EMCS, allowing it to take control of HVAC operational inputs without additional sensors or hardware.

According to the vendor, the software collects numerous data points over time including occupancy schedules, HVAC or whole building power consumption, and outdoor air temperature. The software uses these inputs to develop a multi-variable regression model of the building's thermal loads, as well as a baseline model of the building that predicts HVAC power consumption for each hour of the year. Based on the thermal model, the software tests various HVAC control strategies for the upcoming 24 hour period in order to find the control approach that saves the most energy compared to the baseline model. The software continually updates the building thermal model every two hours and further refines the HVAC control approach based on the updated real-time data.

The following examples illustrate how the software uses predictive models to optimize HVAC operations:

- The building space temperature is modulated within a predetermined comfort range; usually between 70 °F and 74 °F. By modulating space temperature, rather than operating at a fixed set point, the software can shift HVAC power consumption to times of day when compressor efficiency is greater and energy costs are lower. On a day that the software predicts will have a warm afternoon, the building is pre-cooled in the morning and internal temperatures are allowed to 'float' upward in the afternoon. In theory, this results in an overall reduction in HVAC energy consumption because the HVAC systems operate more efficiently during the cool morning than they would in the midday heat. Additionally, the load is shifted away from peak hours when energy prices and grid demand are highest.
- The software optimizes economizer operation in a similar fashion. If the software predicts that a warm afternoon is approaching, but free cooling is available in the morning, then the software will maximize the free cooling capability by dropping the space temperature set point in the morning and brings in as much cool outdoor air as possible to pre-cool the space. Alternatively, if the software predicts a cooler afternoon approaching, then it limits free cooling in the morning, maintains a higher space temperature, and limits the reheating required to meet the afternoon heating load.
- The software communicates directly with utility AutoDR servers[1] to automatically enact demand response measures that are optimized based on the predicted building thermal conditions. Using the predictive model, the software tests numerous demand response scenarios in order to find the approach that will yield the highest demand reduction without impacting occupant comfort. A standard EMCS typically responds to a DR signal by adjusting HVAC set points using assumed temperature and load profiles based on a

[1] To participate in California's Automated Demand Response (AutoDR) incentive programs, building owners must employ controls that automatically initiate pre-programmed demand response strategies in response to signals sent from utility servers.

design-day model of the building. The predictive software, on the other hand, uses real-time data to fine-tune the scale and timing of the temperature adjustments based on the specific space cooling loads of that day.

In addition to the control functions performed automatically by the software, the software vendor provides on-call support to facilities staff and some basic remote fault detection capabilities. This component of the offering is as crucial as the energy-saving algorithms. Facility operators often do not like outside sources controlling their systems, which is specifically what these software packages are designed to do. Without communication channels, facilities staff would likely disable or bypass the system if an issue arose, rather than working through the issue via on-call support. Therefore, having direct communication between the software vendor and facilities staff is an important component in ensuring that whatever savings the software delivers are able to persist over the long term.

Case Study

To test the vendor's savings claims, as described in the section above, kW Engineering conducted a two-year M&V study from May 2013 to August 2015 in a large office building in San Diego, California.

Baseline Building Description

The test building was 6 floors, 144,000 conditioned square feet, and was built in 2001. The HVAC equipment included three 65-ton and three 72-ton Trane Intellipak packaged VAV air handling units. Each unit was equipped with a VFD-controlled supply fan, water-side economizers, water-cooled DX compressors for cooling, and hot water coils for pre-heating. The DX compressors were on a common condenser water loop served by a rooftop cooling tower that operates with a VFD-controlled fan. Hot water was generated by a rooftop boiler that used VFD-controlled pumps to deliver hot water to the air handling units and to perimeter VAV boxes for additional zone-level heating. Prior to installing the predictive HVAC controls, the building's JCI Metasys DDC system had full control of all HVAC operations and set points. This EMCS controlled HVAC scheduling, modulated fan and pump VFD speeds to maintain constant pressure set points, operated a supply air temperature (SAT) reset based on internal building loads, staged compressors to maintain SAT set point, and implemented water-side economization. The control strategies operating in this building are consistent with many other large office buildings throughout the state, and meet many of California's Title 24-2013 new building design standards.

Project Goals

kW Engineering and SDG&E designed the M&V case study to quantify the impact of predictive HVAC control software in the following areas of interest:

- Average annual reduction in the building's electrical energy consumption during a typical meteorological year (TMY).
- Demand reduction potential and ability to communicate automatically with utility demand response automation servers (DRAS).

- Impact on occupant comfort during a demand response event and throughout a typical year.
- Interaction with on-site facilities staff.

We established these areas of interest to directly test the vendor's energy savings claims, to investigate the non-energy impacts of the software, and to determine the potential benefits and barriers to adopting the technology on a wider scale. The following sections describe the approach used to test the equipment, the results of the testing, and the conclusions that can be drawn from the assessment.

Testing Approach

To test the impacts of the predictive control software, kW Engineering conducted an M&V study that applied the International Performance Measurement and Verification Protocol (IPMVP) "Option B: Retrofit Isolation" approach. This approach required measurement of all sub-systems within the building that were directly affected by the investigated technology, in order to confirm the energy savings within a reasonable margin of error. We applied this M&V approach to the areas of interest listed above as follows:

Average Annual Electrical Energy Reduction

To quantify the average annual reduction in electric consumption, we monitored the power of all HVAC systems (air-side and water-side) in 15-minute intervals for 18 months prior to the software installation and 10 months following the installation. We collected these power readings from the utility-grade meter that is dedicated solely to the HVAC equipment throughout the building. We measured power and current from individual sub-systems as well to confirm that the meter-level power readings were reasonable.

We used the collected power data, outdoor air temperature data from a local NOAA weather station, and building occupancy profiles to develop a multi-variable regression analysis that models the baseline and post-retrofit annual HVAC power consumption. To develop the regression, we used the statistical computing software 'R,' and applied the regression technique described in LBNL-4944E, an April 2011 article from Lawrence Berkeley National Labs that outlines recommended procedures for quantifying changes in building electricity use (Mathieu et al., 2011). To apply this calculation approach, which is too extensive to describe in full herein, our first step was to model building electric load as a function of outdoor air temperature, occupancy status, and time-of-week. To develop this portion of the model, we divided the week into 60-minute time intervals (other intervals can be used). For each interval, we calculated an individual baseload. This baseload depends only on time, not on outdoor air temperature or occupancy. In addition to the baseload, we calculated several temperature-based load components. To determine the temperature-based load, we applied a simple linear outside air temperature function to model the building's unoccupied operation; and a piecewise linear outside air temperature function using six equally-sized segments to model the temperature-based load while the building is occupied. To determine the total load for a given time, we summed the temperature-based loads, occupancy-dependent loads, and the baseload for all given time, occupancy, and weather conditions. As a result there are 175 components calculated during the regression: 168 for the time-of-week factor, 6 for the occupied piecewise linear outside air temperature function, and 1 for the unoccupied temperature function.

To calculate energy savings over a full typical year, we applied local TMY-3 weather to both the baseline and post-retrofit regression models in order to calculate yearly, adjusted baseline and proposed demand profiles.

The regression model's accuracy depended on having sufficient data to cover all outdoor air temperature ranges, so long-term monitoring was required. It was also imperative that no changes to the building operations, other than the change being measured (in this case, the implementation of a cloud-based control software), occurred between the baseline and post-retrofit data collection periods. This was essential in order to provide an 'apples-to-apples' comparison of the baseline and post-retrofit data. To this end, we culled the data set to eliminate any baseline data that was not representative of the building's loads throughout the monitoring period. For example, the building's facilities staff informed us that floor #5 became unoccupied at the end of 2013, during our baseline monitoring period, and continued to be unoccupied through the post-retrofit M&V period. Therefore, we removed all 2013 data from the baseline analysis to provide a like-for-like comparison between the baseline and post-retrofit cases.

We assessed the fit of our baseline and post-retrofit models based on the R-squared value and the coefficient of variation of the root mean squared error [CV(RMSE)], as listed below.

Baseline Regression Model: CV(RSME) = 18.3%; R-squared = 0.954
Post-Retrofit Regression Model: CV(RSME) = 19.5%; R-squared = 0.952

Figure 1, below, graphically illustrates how accurately the regression model predicts baseline power consumption both during the occupied and unoccupied periods.

Figure 1. Data sample from baseline trending period comparing measured HVAC power consumption over time (blue line) compared to the regression model's predicted power consumption (orange line)

The R-squared, CV(RSME), and graphical examples indicate that the baseline and post-retrofit models were a reasonable fit to the collected monitoring data. Since we calculated energy savings as the difference between the baseline and post-retrofit energy use, the uncertainty of the savings claim is affected by the accuracy of both the baseline and post-retrofit models. To validate the savings claims, we compared the uncertainty of the calculated savings to the widely-used standards in ASHRAE Guideline 14. This uncertainty assessment is presented in the Test Results section below.

Demand Reduction Potential

To measure the demand reduction, we conducted a 2-hour demand response test and a 4-hour demand response test, during which we simulated a DR event signal from SDG&E's DRAS server[2], then tracked the resulting change in building demand. To measure the savings, we first established a baseline demand profile using a 10-day average baseline with "day-of" adjustment. This approach was consistent with the method that the California Investor Owned Utilities (IOUs) use to calculate savings for their DR incentive programs, and has been vetted by the California Public Utilities Commission (SCE, 2013). This method required 15-minute interval data from the HVAC meter to determine the hourly average load profile from the 10 business days prior to each demand response event. Using this interval data we calculated the average baseline demand for each hour, then applied the "day-of" adjustment ratio. To determine this ratio, we calculated the average kWh usage of the first three of the four hours before each DR event, and divided this value by the average kWh usage for the same three hours from the 10 days prior to the event. This adjustment factor was designed to account for changes in weather patterns or building operating conditions during the day of the event compared to the previous 10 days. The result was a baseline load profile that predicts the hour-by-hour HVAC power consumption for the demand response day, without any DR procedures in place. The difference between the predicted power consumption and the actual measured HVAC consumption measured during the DR event is equivalent to the demand response load shed.

Building Occupant Comfort Assessment

Though the exact control methodology of the studied software is proprietary, the vendor stated that their control logic makes set point adjustments that affect both the supply air temperature and airflows in order to achieve both energy and demand response savings. Despite adjusting these set points, the manufacturer claims that their software does not significantly impact occupant comfort. We identified occupant comfort as a critical aspect of this case study, because in order for an HVAC control technology to provide persistent energy savings over time (i.e., is not decommissioned, or overridden by site staff due to occupant complaints), the occupants must remain comfortable.

For the purposes of this study, we defined 'occupant comfort' quantitatively as maintaining space temperature and relative humidity within acceptable ranges, as defined by ASHRAE Standard 55-2013. To determine if the internal conditions were comfortable, we measured space temperature and relative humidity inside one open office area on each floor during both the pre-retrofit and post-retrofit monitoring periods. We placed sensors away from exterior walls and from supply air registers in order to track average conditions in the room as closely as possible. We also assessed occupant comfort qualitatively by asking the building's facilities director to track hot or cold calls from the tenants and identify if there was any marked increase in tenant complaints after the software was implemented.

[2] The Demand Response Automation Server (DRAS) automatically sends a signal to the building on demand response days in order to initiate the building's pre-programmed DR control sequences.

Facilities Staff Interaction

To determine the facilities staff's overall opinion of the studied software and the remote management services the vendor provided, we interviewed facilities staff regarding and how they interacted with the software on a daily basis. Interview questions primarily sought to determine whether the software was a burden or benefit to on-site staff and their day-to-day workload.

Test Results

The result of the multi-year M&V investigation revealed that the software did achieve verifiable annual electrical energy savings, was capable of implementing automated demand response to achieve meter-level demand shed, and achieved both of these goals without significantly impacting occupant comfort. The following sections summarize the data and analyses that led to these results.

Annual Electricity Reduction Results

Based on the results of the multi-variable regression analysis, we determined that the building's average annual HVAC electrical energy consumption dropped by 10.7% as a direct result of the new cloud-based EMCS controls. The project goals required that the uncertainty meet ASHRAE Guideline 14 standards, which states that the uncertainty shall be less than 50% of the reported annual savings (ASHRAE, 2014). We calculated an uncertainty of 4.4% at 68% confidence, which per ASHRAE Guideline 14 standards, indicated that the savings estimates were within a reasonable margin of error. Figure 2 illustrates the average reduction in total HVAC energy consumption across all observed load ranges during the two-year testing period.

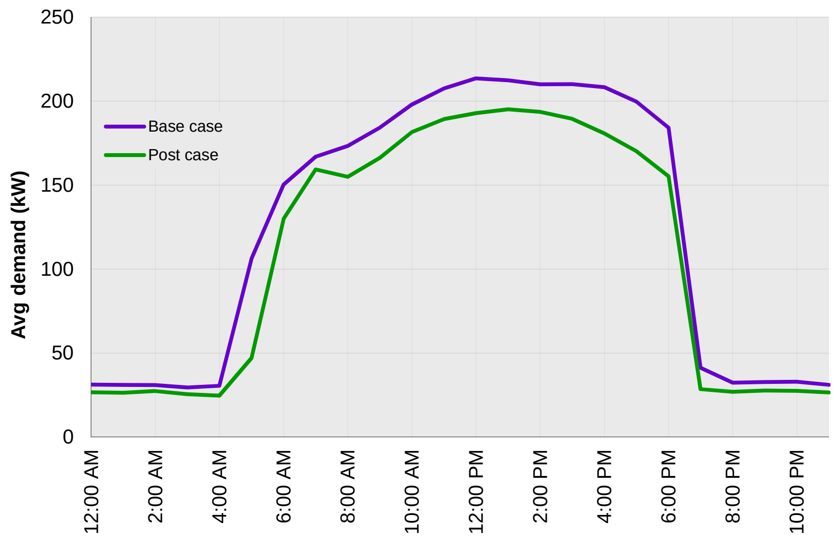

Figure 2. Graph of baseline and post-retrofit average weekday HVAC power consumption data

Demand Response Results

The 2-hour DR test yielded an average 8% reduction in HVAC equipment demand, while the 4-hour DR test yielded an average 4% demand reduction. Figure 3 illustrates the average hourly HVAC power consumption during the 4-hour DR test, as well as the unadjusted and adjusted baselines used in the 10-day average baseline with "day-of" adjustment demand reduction analysis.

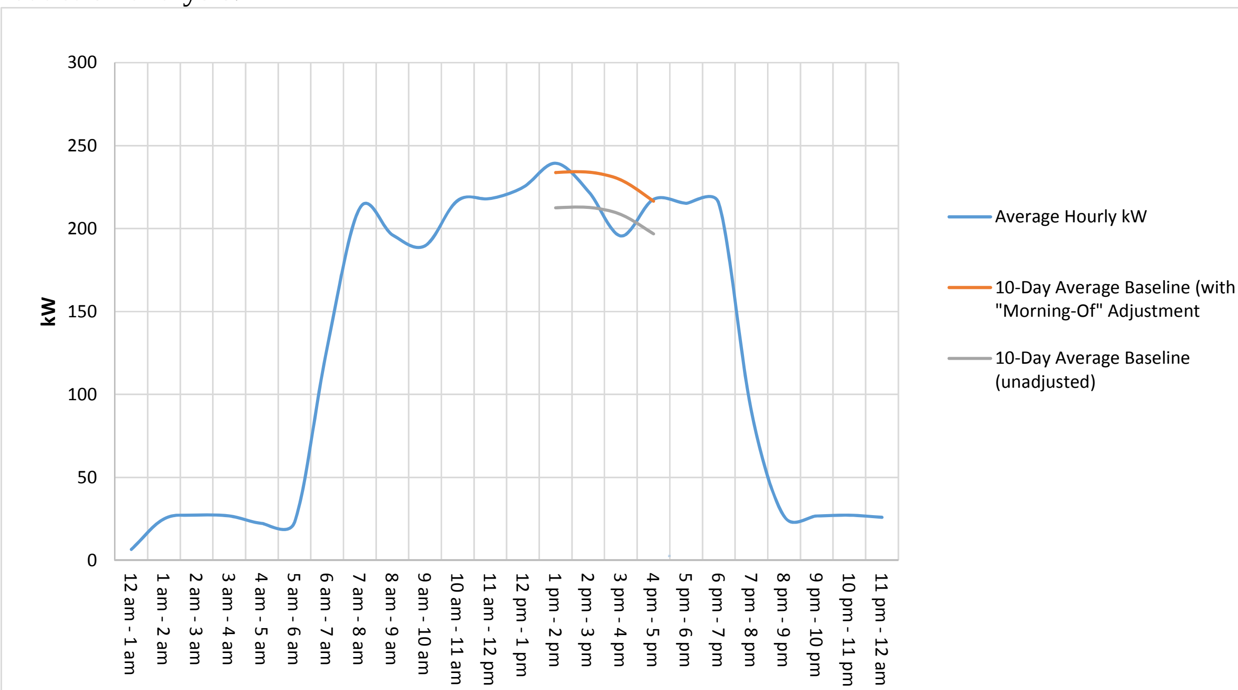

Figure 3. Measured data and baseline data from 4-hour DR test, conducted on July 22, 2015

The graph shows a clear demand reduction from the adjusted baseline, but also shows that the reduction varies significantly throughout the demand response period.

Occupant Comfort Results

We compiled and graphically analyzed the space temperature and relative humidity data that we collected from each floor. The data showed that baseline and post-retrofit space temperature varied from -0.6 °F to +1.3 °F of the space temperature set point. The accuracy of the temperature sensors was +/- 0.63 °F, and therefore the total uncertainty in the temperature measurements was +/- 1.26 °F. The maximum temperature increase we observed was only 0.04 °F greater than the measurement uncertainty, indicating no significant change in space temperature after the software was installed. Additionally, we found the humidity ratio to be constant at 0.010 +/- 0.001 throughout the monitoring period. ASHRAE Standard 55-2013 states that the temperature can range from 67 °F to 82 °F and humidity ratio must be maintained at or below 0.012 to maintain occupant comfort (ASHRAE, 2013).

Figure 4 and Figure 5 show the average indoor space temperature and humidity readings throughout both the baseline and post-retrofit monitoring periods.

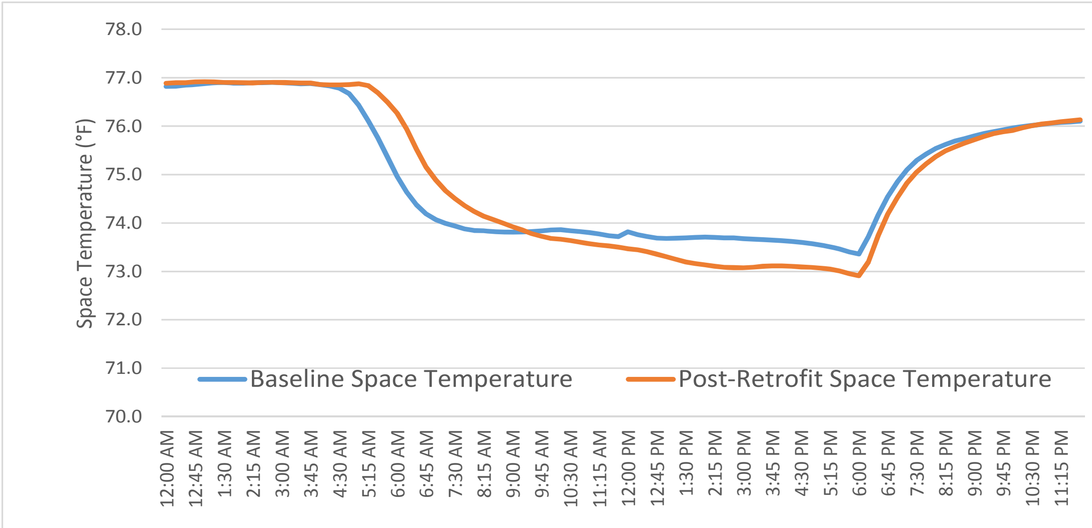

Figure 4. Average of weekday interior space temperatures from baseline and post-retrofit monitoring periods

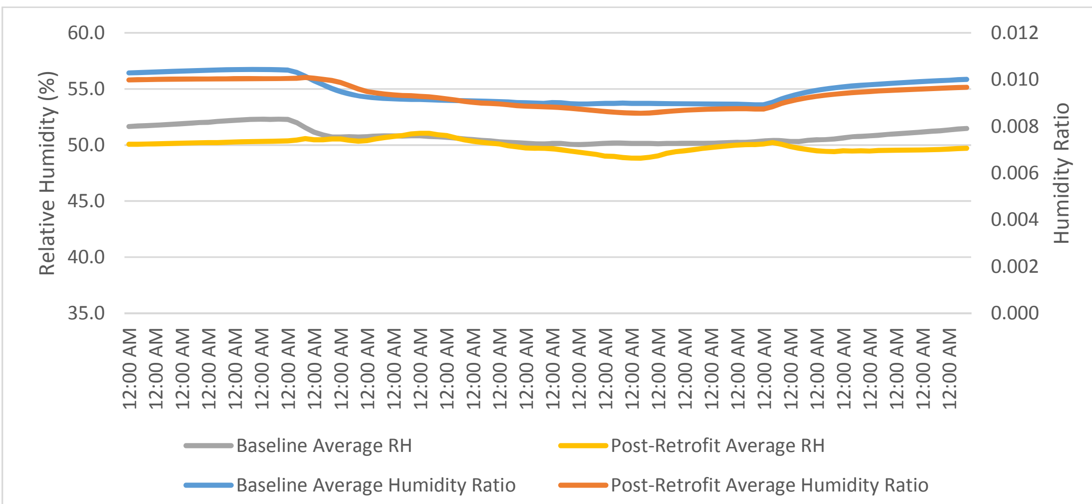

Figure 5. Average of monitored weekday relative humidity and humidity ratio from baseline and post-retrofit monitoring periods

The data also indicated that the building remained within the comfortable temperature and humidity ranges throughout the DR testing as well. Furthermore, the on-site facilities staff who field tenant hot and cold calls confirmed that there had been no increase in the frequency of tenant complaints after implementation of the software.

Facilities Staff Interaction Results

During the interviews we conducted with facilities staff after the post-retrofit M&V period, their overall opinion of the software and the services provided by the vendor was high. No staff indicated any increase in their workload due to implementation of the software. The director of facilities stated that he had used the on-call technical support service on a few occasions to report hot or cold calls from his tenants, and that the software provider addressed these issues within a reasonable timeframe. He noted that the software vendor had alerted him of a number of potential operational issues using their fault detection as well. Furthermore, at no point during the testing period did facilities staff disable the predictive control system due to operational concerns, though some internet connectivity issues did result in communication problems with the cloud server at times. Overall, the facilities staff agreed that the software did not impinge on their ability to work in an efficient manner, and stated that overall they found it an improvement in the operation of their building.

Cost-Effectiveness

To assess the market readiness for widespread implementation of the predictive controls, we calculated the project payback and compared the payback period to the expected useful life of the software. In this particular application, we calculated that the customer's simple payback was 6.5 years before incentives, and 4.8 years with assumed utility incentives (calculated assuming the technology was incentivized through existing SDG&E customized incentive programs). The typical contract period for the software is 5 years, after which point additional payments to the software vendor would be required. Therefore, for this particular building, the utility incentive was necessary in order to bring the payback within the terms of the initial service agreement (SDG&E, 2015). These incentive programs, which had a state-wide budget of $1.7 billion in 2014 (DOE: 2015), provide significant funding for customers to install new, cutting-edge technologies. As such, these programs are major drivers for change in the industry, and push the state closer to its lofty efficiency goals. Our case study confirmed that this software could achieve verifiable energy savings, and the simple payback assessment indicated that utility incentives could be critical to bringing the project cost-effectiveness in line with customer goals. However, the long-term M&V required to verify the savings still remains a potential barrier to incorporating these technologies into custom incentive programs. Until then, pay-for-performance programs are likely the most effective method of incentivizing this technology.

Conclusions

The results of our case study conclusively showed that this particular cloud-based predictive HVAC control software successfully reduced the building's annual HVAC electricity consumption within the vendor's predicted range, reduced peak demand in response to a DR signal from the utility's automated server, and achieved these goals without significantly impacting occupant comfort. These results are significant as they represent the first independent, data-based verification of one such technology's saving potential.

These results are not enough, however, to make broad-ranging claims about the technology's savings potential in other applications. The software's energy savings and demand response potential could vary dramatically based on building operating profile, HVAC system types, building types, and climate zones. For example, this technology relies on predictions of upcoming building operations in order to determine the most efficient HVAC control approach. Therefore, the software may yield less favorable energy savings in a building with non-standard or sporadic occupancy profiles, as the load is harder to predict. Conversely, the software could achieve greater energy savings than we observed in this project if the existing building systems had differed. For example, the investigated software's control algorithms claim to optimize air-side economizer functions, but the studied building was only operating water-side economizers. Therefore, no savings from air-side economizer control improvements could be realized.

Furthermore, we do not have enough data from this initial case study to determine whether one company's predictive software is favorable to another's. When the power conditioning 'black box' devices became popular in the early 2000's, some of the technologies yielded legitimate energy savings, while others were ineffective or even detrimental to a building's power consumption. The same may be true of the different cloud-based predictive HVAC control technologies that are emerging on the market today.

To evaluate the performance of this software across a wider range of applications, and to test the capabilities of multiple software vendors, we recommend that vendors, utilities, or other interested parties commission additional independent case studies. The data from these additional case studies could then be aggregated to identify trends in what building types yield the greatest savings, and which particular software providers yield the most reliable, positive results.

Despite the above caveats, we know from this test that the software achieved energy savings beyond what the existing EMCS had provided, that these savings persisted over time, and that we could conclusively measure the savings. We believe that this initial study demonstrates the potential for cloud-based predictive HVAC controls to significantly improve HVAC efficiency in buildings with existing EMCS systems. Once fully tested, vetted, and accepted on a more widespread basis, we expect that these predictive controls will play a major role in driving the state toward its 2030 efficiency goals and beyond.

References

American Society of Heating Refrigeration and Air Conditioning Engineers (ASHRAE). 2013. "Thermal Environmental Conditions for Human Occupancy." *Standard 55-13.* ASHRAE.

ASHRAE. 2014. "Measurement of Energy and Demand Savings." *Guideline 14-2014*. ASHRAE.

CEC. 2012. "2013 Building Energy Efficiency Standards for Residential and Nonresidential Buildings." Title 24, Part 6. CEC-400-2012-003-CMF-REV2. California Energy Commission. http://www.energy.ca.gov/2012publications/CEC-400-2012-004/CEC-400-2012-004-CMF-REV2.pdf

Itron. 2006. "California Commercial End-Use Survey." CEC-400-2006-005 [pg 158]. California Energy Commission

Mathiew, J., P. Price, S. Kiliccote, and M. Piette. 2011. "Quantifying Changes in Building Electricity Use, with Application to Demand Response." *LBNL-4944E.* Ernest Orlando Lawrence Berkeley National Laboratory. http://eande.lbl.gov/sites/all/files/LBNL-4944E.pdf

Pacific Gas and Electric Company (PG&E). 2004. "Before you Buy a Black Box." *Power Quality Bulletin No. 1.* PG&E. http://www.pge.com/includes/docs/pdfs/about/news/outagestatus/powerquality/black_box.pdf.

Pande, A., and C. Austin. 2012. "Estimation of EMCS Presence in Commercial Buildings in PG&E Territory and a Snapshot of Technologies in the EMIS Landscape." *ET Project Number: ET11PGE4221* [page 12]. Heschong Mahone Group, Inc.

SB 350, De Leon. "Clean Energy and Pollution Reduction Act of 2015." 2015. https://leginfo.legislature.ca.gov/faces/billNavClient.xhtml?bill_id=201520160SB350

SCE. 2013. "10-Day Average Baseline and "Day-Of" Adjustment." NR-2225-V1-0413. Southern California Edison. https://www.sce.com/wps/wcm/connect/c47fe131-45a5-4456-b3ad-f8a781a42bd6/10+Day+Avg+Baseline+%26+Day+Of+Adj_+NR-2225-V1-0413.pdf?MOD=AJPERES.

SDG&E. 2015. "Model-Based Predictive HVAC Control Enhancement Software." *Emerging Technologies Coordinating Council*, DR13SDGE0006 Report. San Diego Gas & Electric.

US Department of Energy. 2015. *Energy Incentive Programs, California*. Washington, DC: United States Department of Energy (DOE). http://energy.gov/eere/femp/energy-incentive-programs-california.