Using Models to Provide Predicted Ranges for Building-Human Interfaces

Preprint

Nicholas Long, Jennifer Scheib, Shanti Pless, and Marjorie Schott
National Renewable Energy Laboratory

Presented at the 2013 Building Simulation Conference
Chambéry, France
August 25-28, 2013

NREL is a national laboratory of the U.S. Department of Energy
Office of Energy Efficiency & Renewable Energy
Operated by the Alliance for Sustainable Energy, LLC

Conference Paper
NREL/CP-5500-58521
September 2013

Contract No. DE-AC36-08GO28308

NOTICE

The submitted manuscript has been offered by an employee of the Alliance for Sustainable Energy, LLC (Alliance), a contractor of the US Government under Contract No. DE-AC36-08GO28308. Accordingly, the US Government and Alliance retain a nonexclusive royalty-free license to publish or reproduce the published form of this contribution, or allow others to do so, for US Government purposes.

This report was prepared as an account of work sponsored by an agency of the United States government. Neither the United States government nor any agency thereof, nor any of their employees, makes any warranty, express or implied, or assumes any legal liability or responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represents that its use would not infringe privately owned rights. Reference herein to any specific commercial product, process, or service by trade name, trademark, manufacturer, or otherwise does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States government or any agency thereof. The views and opinions of authors expressed herein do not necessarily state or reflect those of the United States government or any agency thereof.

This report is available at no cost from the National Renewable Energy Laboratory (NREL) at www.nrel.gov/publications.

Available electronically at http://www.osti.gov/bridge

Available for a processing fee to U.S. Department of Energy and its contractors, in paper, from:

U.S. Department of Energy
Office of Scientific and Technical Information
P.O. Box 62
Oak Ridge, TN 37831-0062
phone: 865.576.8401
fax: 865.576.5728
email: mailto:reports@adonis.osti.gov

Available for sale to the public, in paper, from:

U.S. Department of Commerce
National Technical Information Service
5285 Port Royal Road
Springfield, VA 22161
phone: 800.553.6847
fax: 703.605.6900
email: orders@ntis.fedworld.gov
online ordering: http://www.ntis.gov/help/ordermethods.aspx

Cover Photos: (left to right) photo by Pat Corkery, NREL 16416, photo from SunEdison, NREL 17423, photo by Pat Corkery, NREL 16560, photo by Dennis Schroeder, NREL 17613, photo by Dean Armstrong, NREL 17436, photo by Pat Corkery, NREL 17721.

Printed on paper containing at least 50% wastepaper, including 10% post consumer waste.

USING MODELS TO PROVIDE PREDICTED RANGES FOR BUILDING-HUMAN INTERFACES

Nicholas Long, Jennifer Scheib, Shanti Pless, and Marjorie Schott
National Renewable Energy Laboratory, Golden, CO, USA

ABSTRACT

Most building energy consumption dashboards provide a snapshot of building performance; others provide detailed historical data for comparison to current usage. This paper discusses the Building Agent platform, which was developed and deployed in a campus setting at the National Renewable Energy Laboratory as part of an effort to maintain the aggressive energy performance achieved in newly constructed office buildings and laboratories.

The Building Agent provides aggregated and coherent access to building data, including electric energy, thermal energy, temperatures, humidity, and lighting levels, and occupant comfort feedback, which are displayed in various manners for visitors, occupants, facility managers, and researchers. This paper focuses on the development of visualisations for facility managers, or an energy performance assurance provider, where metered data are used to generate models that provide live predicted ranges of building performance by end use. These ranges provide simple, visual contexts for displayed performance data without requiring users to assess historical information or trends. Several energy modelling techniques were explo red, including static lookup-based performance targets, reduced-order models derived from historical data using main effect variables such as solar radiance for lighting performance, and integrated energy models using a whole-building energy simulation program. This paper describes the new building construction backdrop that has motivated this work, the system architecture providing access to building data, the various modelling approaches currently employed, and the visualisation methods used to display performance and modelled data.

INTRODUCTION

As net-zero energy building procurement moves out of infancy, the endeavour expands from realizing aggressive energy goals in design to maintaining designed performance over the life of the building. Changes in occupancy patterns, miscellaneous loads, and installed equipment performance can significantly impact year-to-year energy performance. In one building lifetime scenario, energy use increases over time as equipment is added or operational schedules change. In another, energy use decreases over time as equipment degrades or fails but at the sacrifice of occupant comfort. The ideal scenario is when equipment is maintained and upgraded appropriately, which reduces energy consumption and maintains occupant comfort continually (accounting for variations in weather and other external factors). To ensure the preferred scenario, an interface between the building, the occupants, and the operational staff is needed for fact-based, proactive decision-making.

Like energy models in design, dashboards have surfaced as the primary interface between the building and the decision makers. Many dashboard visualisations are paired with a single energy goal, such as net-zero energy, which is sufficient if no problems are encountered but does not provide direction for corrective action if the goal is not met. The premise of this paper, then, is that the dashboard must not oversimplify the design-based energy model results but rather disentangle them for ongoing analysis by building decision makers. The result is a more complex dashboard architecture and a deeper understanding of the building's operation; both outcomes will be illustrated through the U.S. Department of Energy's (DOE) National Renewable Energy Laboratory (NREL) campus dashboard and occupant interface, the Building Agent (BA).

NREL Campus Background

The DOE/NREL goal for the campus is to expand its leadership as a state-of-the-art laboratory that supports innovative research, development, and commercialization of renewable energy and energy efficiency technologies that address the nation's energy and environmental needs. This growth has resulted in a significant increase in employees and facilities on its 327-acre main campus in Golden, Colorado.

To support this growth over the last five years, NREL developed and demonstrated new construction procurement methods that proved cost effective and showed that 50% energy savings versus typical building code are possible when design-build teams integrate to achieve specific and measurable whole-building energy goals. NREL facility growth provided an opportunity to demonstrate the integrated approach in real projects by incorporating energy performance specifications into the projects' contracts. NREL developed and piloted this energy performance-based design-build process in 2008 with the Research Support Facility (RSF I), an 824-occupant, 220,000-ft^2 office building. The process has since been replicated in other campus projects such as an office wing expansion, the Research Support Facility II (RSF II), a smart grid laboratory, Energy Systems Integration Facility, a parking structure, and a site entrance building.

Each project incorporated world-class efficiency strategies using contractual energy use requirements in the design-build contracts, all on typical DOE construction budgets. In addition to general energy reduction goals such as a 50% reduction versus typical building code and sustainability goals such as Leadership in Energy and Environmental Design Platinum certification, the contract included specific and directly measurable energy use requirements:

- RSF I: 36 kBtu/ft^2/yr (11 kWh/m^2/yr) energy use intensity (EUI), including the data centre and net-zero energy
- RSF II: 33 kBtu/ft^2/yr (10 kWh/m^2/yr) EUI and net-zero energy
- ESIF: 27 kBtu/ft^2/yr (9 kWh/m^2/yr) EUI for the office space, and 1.06 power usage effectiveness (PUE) for the data centre
- Parking Structure: 175 kBtu/parking space/yr (51 kWh/parking space/yr)
- Site Entrance Building: 32 kBtu/ft^2/yr (10 kWh/m^2/yr) EUI and net zero energy

Each project's design-build team successfully designed and constructed these buildings to reach the contractual energy goals, substantiated by energy models and submetering during the first year of operation. NREL terms this contractual energy goal substantiation effort, combined with energy goal maintenance over the life of the building, energy performance assurance.

Energy Performance Assurance

In each of NREL's recent campus construction projects, whole-building energy models were started in the pre-design phase by each proposing design-build team to prove that its design concepts were likely to meet the contractual goals (Pless, Torcellini, Shelton, 2011). The design-based energy models evolved over the design phases and were updated based on constructed reality. This energy performance assurance task differs from a typical design process in that information from the design-based energy model was used to make iterative decisions about the building and understand cost and energy performance trade-offs versus simply verifying that design concepts perform as expected at the end of design.

The first step in energy performance assurance is to evolve energy models throughout the design and construction. Specifically, this means continually refining building geometry, weather input, occupancy schedules, etc. For the NREL campus, energy-modelling reports were provided at each phase of design and as energy-cost trade-off questions arose, with clearly identified variables such as input schedules that could be discussed among the integrated owner, design, and construction team.

The second step is submetering. This requires the team to consider the electrical design in order to layout meters appropriately to make end use aggregation available for validating the design-based energy models. For each NREL project, submetering was a contractual requirement beyond the implicit contractual need to verify the energy goal at occupancy.

These first two steps are critical to setting up a building-occupant interface such as a dashboard that allows for on-going, proactive decision-making, as the evolving energy model allows for in design.

The final step in energy performance assurance, and the focus of the remainder of the paper, is the setup of the building-occupant interface. This effort includes defining the system architecture for data collection and designing system visualisations that allow occupants—including facility managers—to make goal-driven decisions.

Building Agent is NREL's building-occupant interface, developed in tandem with the recent campus construction. The following BA description focuses on the RSF because the building has been operational for more than one year, providing a fuller picture of the effort required to maintain energy performance. The RSF was the living laboratory for BA development.

SYSTEM DESIGN

The BA scope extends beyond a typical energy dashboard; it collects and displays energy performance data and allows for analysis and investigation into the balance of energy use, energy costs, and comfort by facility managers and typical occupants.

Building Agent Architecture

The BA architecture consists of four layers: hardware and protocols, databases, applications, and visualisations. *Figure 1* defines the layer components and colour codes them according to those developed uniquely for BA ("Building Agent" and "Visualisations") and a shared campus resource ("DataBUS").

The first layer addresses the spatially distributed and varied protocols for monitoring energy- and comfort-related events on campus. Elements of this layer include (but are not limited to):

- Power meters such as lighting, mechanical, miscellaneous loads, elevators, and photovoltaic (PV) panels (*Modbus* protocol) on each building subsystem
- Flow meters on hot and chilled water (*BACnet* protocol)
- Spatial *environmental sensors* such as illuminance, temperature, and humidity sensors at individual workstations
- *Other protocols* such as weather station sensors, including global horizontal and vertical irradiance, dry bulb temperature, and wind speed and direction
- *Client applications* deployed to collect occupant feedback. Occupants are considered meters with many sensors in the BA architecture.

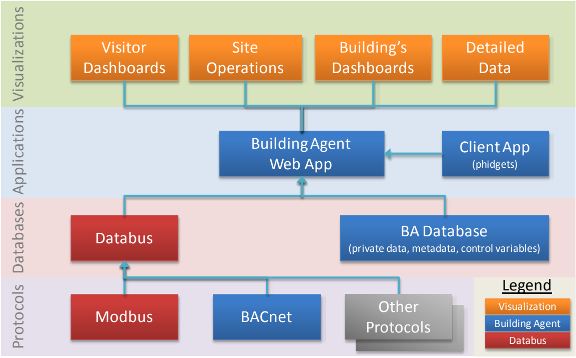

Figure 1: BA Architecture

As a second layer, all the devices either send data or are polled at a time interval appropriate to the rate of change of the condition being measured. For example, environmental sensors send information when occupants at workstations give feedback about their comfort or every 10 minutes. The power meters are set to send data at one-minute intervals. All these data are collected, cleansed, aggregated, and tagged in two databases for use by the BA applications. One database, DataBUS, houses raw building performance data. These data are open and available for a variety of NREL research and assessment needs. A second database, specific to and named for BA, houses metadata for the raw DataBUS data used for campus dashboards. The BA database also segregates private occupant data such as comfort feedback.

The third and fourth layers of the BA architecture go hand-in-hand. The data are meaningfully organized and manipulated and then displayed via a series of dashboards. The up-to-date status report from the building to the occupant allows for a human-in-the-loop control scenario. For the RSF, the interface takes the form of data visualisations through a desktop client application, a campus display, and a website. Each instance is meant to connect with occupants in different scenarios but all communicate the RSF's current performance and clearly indicate when the building is not meeting expectations.

The design team was contractually required to deliver a dashboard for the RSF (see *Figure 2*). It used the submetered data from the power quality and flow meters, but the data analysis layer was not included in the architecture. This means that there was no context for the displayed energy performance. Further research into historical information was necessary for performance assurance tasks such as verifying the building's contractual energy goal was being met or determining the appropriate course of action if the building was not meeting the goal.

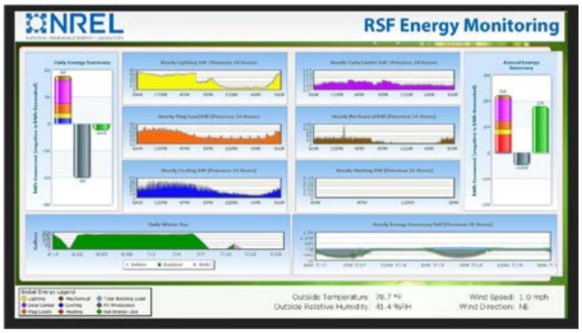

Figure 2: Original, Control Vendor-Provided Dashboard for the RSF (July 18, 2011, sunny day)

The second iteration of the RSF dashboard design added the data analysis layer. Over the first year of building operations, the metered data were collected and compared to the energy model results at the whole-building and end-use scales, as is done in an enhanced measurement and verification scope (Sheppy 2013). Tweaks were made to building systems and meters as needed to ensure that the RSF's energy performance met the contractual energy goal. A new interactive dashboard (see *Figure 3*) was developed that displayed the building's power profile as a breakdown of the whole-building energy goals.

Figure 3: BA Dashboard for the RSF (whole building, May 15, 2013, partly cloudy day)

The interface enables users to "drill down" to inspect individual end use time series data as shown in *Figure 4*.

Figure 4: BA Dashboard for the RSF (data centre power consumption (kW), February 13, 2013, energy use not within predicted range)

These examples use power and energy readings to display different aspects of performance. Both are necessary depending on the metric for comparison; multiple dashboard views provide the diversity of information occupants and building engineers need to make informed decisions about comfort and energy use. Although the fourth layer (building-occupant interface, display selection and design) is an important discussion point for meaningful and actionable visualisations (Schott et al., 2012), the remainder of this paper addresses the third, data analysis layer.

Predicted Range Approaches

A unique feature of the RSF dashboard, shown in both *Figure 3* and *Figure 4*, is the display of dynamic, acceptable ranges around subsystem goals. Subsystem goals are important for creating a useful dashboard because it is rare that at any given point in time, the annual energy use or power profile will match the exact goals. Three options and the associated steps for determining the predicted ranges of whole-building and building subsystems are:

1. Lookup tables (no asset or process model available)

Lookup tables can be used for energy goal comparison when no model data are available.

- Start with the whole-building energy goal or first year aggregation
- Use public energy end use breakdown information (DOE, 2010) to divide the total operational energy use among end uses
- Use case studies, rules of thumb, and design guide support documents (Leach, Lobato, Hirsch, et al. 2010) to determine a range of acceptable energy use

The advantages to using lookup tables are that they can be quickly developed, they do not require initial metered data, and they allow a building to be compared to industry standards. The disadvantages are that a specific system's energy use might be unrealistic to take action to make the buildings end use fall in range. Also, the ranges will be static and will not account for occupancy or weather variation.

2. Reduced-order model (based on asset model)

A reduced-order model provides quick results based on only a few independent parameters such as time-of-day, outside temperature, and irradiance. Typically, the manual development of a simple reduced-order model would include:

- Starting with the first year metered data and calibrated annual energy model results (and run with updated weather)
- Plotting the data sets to determine if the measured data aligns with the modelled data
- Ensuring, through data inspection, that the compared sets are equal in time period, units, and resolutions, and determine the acceptable energy use and power profile ranges. Specifically, determine whether the design predictions are being met. If not, make this the high end of the range. If the predictions are being met, determine whether the operating energy use can be ratcheted down further based on occupant comfort feedback or consistent occupancy profiles. This can be an automated process, but is often an art in current practice.
- Determining the necessary metrics for the dashboard to give occupants a complete picture of the building's performance. One metric should be the whole-building energy goal such as an EUI, but other options include net-zero energy, PUE, lighting power density, or less specific power profiles.
- Using modelled or measured data, whichever is determined to be the appropriate performance range.
- Plotting the data sets against a variety of factors such as exterior temperature and time of day. This can be automated using mutual information techniques.

Lookup tables can be used in this approach as a variation on the reduced-order model, when the most influential parameter can easily be broken down into a small number of discrete steps such as time of day. Lookup tables are also useful when an operational model is unavailable or operations are changing over time, but the energy use depends on complex control algorithms.

The advantage of this approach is that the ranges will be dynamic, true to the building operations, and will prevent unrealistic ranges from being displayed in the visualisation. The disadvantages are that an initial data collection and review period is needed before the dashboard can be deployed, preferably by an energy engineer. This approach is currently implemented in the RSF BA.

3. Real-time simulation (using an operational model)

Perhaps the most sophisticated method of synthesizing dynamic performance predictions utilizes real-time simulation. This requires:

- Creating an operational energy model that is a combined forward (fundamental or reduced-order calculations) and data-driven (e.g., metered data-derived occupant schedules) approach
- Creating schedules that are true to the ranges observed in the operating building
- Collecting and cleansing real-time weather data
- Running energy models to determine the acceptable hourly, daily, and annual end-use energy ranges
- Enabling model parameter optimization for cost and comfort that dynamically change ranges

The disadvantage in this scenario is clear; simulation setup is time intensive (and changes during the building life cycle). But, if the project contract requires an aggressive energy goal, an as-built energy model would likely exist; along with an architecture for data collection that supports dashboards and operational decision-making. The advantages are that the interface is dynamic enough to move toward automatic action taken on the building control system and predictive control for demand response cycles.

RESULTS

This section discusses the reduced-order model approach to end-use energy range predictions currently implemented in the BA, showing the development process, lessons learned, and positive outcomes of using a predictive interface in energy performance assurance in the RSF.

Photovoltaic Panels

More than 2 megawatts of PV panels are installed on or near the RSF to offset the annual energy use to achieve net-zero energy. Renewable energy was the last consideration in the design process; however, after system efficiency measures, it is presented first because of the relative ease in determining a dashboard representation.

An asset model is needed to predict energy use based on the final panel specification and configuration, but an operational model is unnecessary because there are no controlled components and the prediction is the energy production goal. Figure 5 shows that the continuous reduced-order model was determined by first plotting the modelled production against horizontal radiation, a known dominant state. Boundaries are then visually added and fitted with equations that are implemented in the BA's application layer.

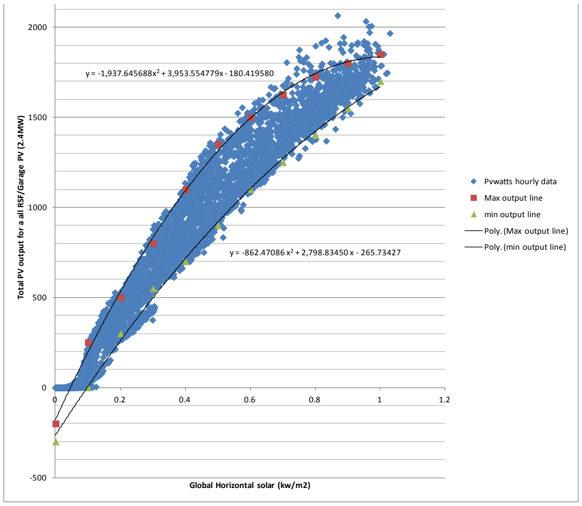

Figure 5: PV Production Versus Global Irradiance

The visualisation layer of the PV predicted range is shown in *Figure 6*.

Figure 6: PV Production Predicted Range Visualisation (May 15, 2013, partly cloudy day)

Lighting

The RSF occupants manually control electric lighting "on" in the RSF; it is automatically controlled "off" by daylight and occupancy sensors and timed sweeps. Daylighting is important in a net-zero energy office building, so a detailed daylighting model was produced in design, the results of which were compared to the first-year measured data as an energy performance assurance step.

Figure 7 shows the importance of an operational model. The nighttime lighting energy use is much higher than predicted in the design model because the system use by a custodial crew was not considered during design.

The annual average daytime lighting energy use matches the model well; therefore, operational data are used to determine the predicted energy ranges.

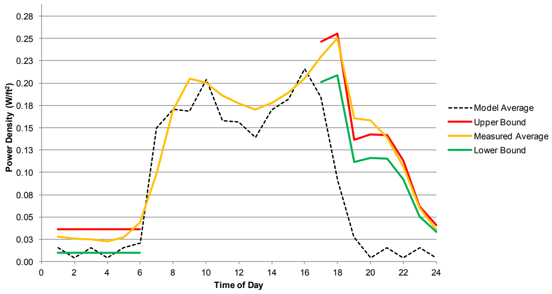

Figure 7: First-year Lighting Operation Hourly Average Power Profile

Figure 8 shows that the daytime lighting energy use depends largely on exterior vertical irradiance. This factor was selected based on experts' judgments because, unlike PV, it was not clear which exterior radiation value is the dominant state. The process for determining the daytime predicted ranges then followed that for PV. The nighttime lighting load is dominated by occupant behaviour that is consistent from day to day, lending itself to a lookup table for the BA application layer.

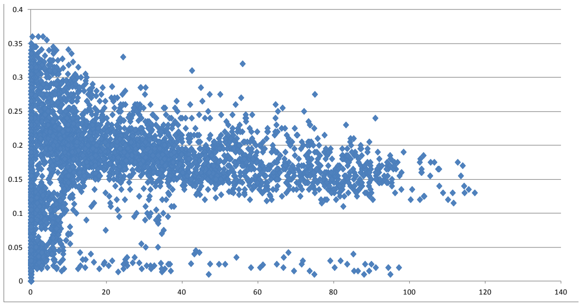

Figure 8: Lighting Power Density (W/m^2) Versus Exterior Vertical Irradiance (W/m^2)

The lighting predicted ranges show a combined continuous reduced-order model and lookup table, as seen in the dashboard implementation in Figure 9.

Figure 9: Lighting Predicted Range Visualisation

Plug Loads

RSF plug loads are regulated; each occupant has a 55-watt allowance for a laptop, monitors, phone, task light, and miscellaneous items. The energy model accounts for this load, but can only infer a diversity factor of occupant use throughout the day.

The first step in operational energy performance assurance is to compare the first-year measured data, or initial operational model, to the design modelled data. Figure 10 shows that the hourly average daytime operational load is much lower than the hourly average design prediction, but the nighttime load is higher. Because the model suggests energy-saving potential, the lower of both datasets at any given hour is used to determine the dashboard energy goal.

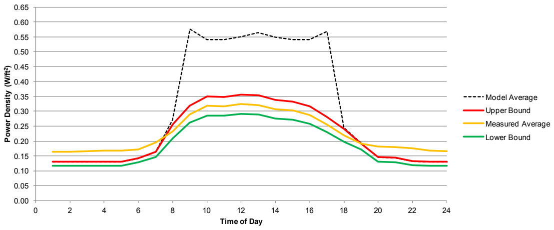

Figure 10: First-Year Plug Load Operation Hourly Average Power Profile

The high and low ranges were determined by visually bounding the first-year data range, excluding outliers, and implementing the range in the dashboard using a lookup table, because time of day is the most obvious factor in the primarily occupant-driven load. The resulting visualisation is shown in Figure 11.

Figure 11: Plug Loads Predicted Range Visualisation (February 13, 2013)

Mechanical

The RSF's mechanical loads consist primarily of ventilation air fans and pumps for radiant heating and cooling fluid. The RSF has hot water and chilled water flow meters to account for its load on the central plant.

The performance assurance and predicted range evaluation for the mechanical systems mirror the process for the plug loads, with two exceptions:

First, the design model predictions overestimate energy use during the day instead of at night, so the first-year metered data are used as an upper bound at night and the design model results are used as an upper bound during the day. This process is demonstrated in Figure 12.

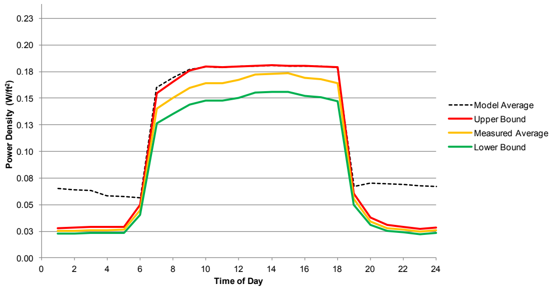

Figure 12: First-Year Mechanical Systems Operation Hourly Average Power Profile

Second, like the plug loads, a lookup table is used in place of a continuous reduced-order model, but this is not the best representation of the performance data, because that hourly and annual variation is much greater than plug loads, which do follow a regular schedule. The mechanical system ranges would be better suited for a real-time energy model in which many factors such as weather, control algorithm, and occupant input can create more dynamic predicted ranges for this end use. *Figure 13* shows the interim result.

Figure 13: Mechanical Systems Predicted Range Visualisation (June 06, 2013)

DISCUSSION

The process used to determine the predicted range models for the RSF, the first building represented by the BA, is presented as an example of the energy assurance process emerging in the commercial buildings industry. The process consisted of design and metered data comparisons, the validation or development of an asset model for systems, and an operational model for control variables, where necessary. Then, a visually determined continuous reduced-order equation or lookup table was applied to the asset and operational models and data, respectively.

For uncontrolled systems such as PV, visually bounding asset model outputs relative to known dominant variables to determine reduced-order equations is a simple and sufficient process. The information that occupants need to know about these types of systems is whether the system is working and whether it is degrading.

On the other end of the spectrum, complex controlled systems for fans and pumps need more rigorous approaches to determining predicted energy use ranges. The reduced-order lookup table approach used for the BA is not sufficient to capture the hourly and daily differences between potential and actual performance, leading to a missed opportunity for energy savings.

As the focus on energy performance assurance sharpens, the commercial buildings industry will develop more rigorous and automated approaches to determining predicted ranges and automatically refining those ranges based on real-time operational models. The RSF process can serve as an example and dataset for that effort.

Value in predicted range visualisation has been realized at NREL through anecdotes to date:

- The high plug load use at night relative to the predicted ranges has lead NREL energy engineers, information technology department, and managers to work together to communicate and give options to employees for turning off monitors, computers, and other miscellaneous loads at night. Best practices were developed based on the lessons learned (Lobato, Pless, Sheppy, et al. 2011).
- A high nighttime lighting load has led to training sessions, reminding custodial crews to turn of ambient lights when they leave an area and to use the egress switches when possible, which often provide sufficient light for the tasks being performed.
- A consistently high data centre load relative to the model gave the data centre manager evidence to request funding for improved hot aisle containment strategies.
- Last but not least, the whole-building energy display comparing the RSF energy use to the contractual EUI and net-zero energy goals has lead to more careful consideration during building upgrade projects. For example, the addition of a 24-hour visualisation room for security was not originally designed for isolated heating, cooling, or ventilation. Isolated systems were added once the building energy engineer was able to show that the building would tip past its annual operating goal if the building systems were used for this system, because the thermal

zones are much larger than the centre being added.

The visualisations have provided information and the ranges have added the needed justification for building energy projects.

CONCLUSIONS

The RSF dashboard development case study reveals the importance of defining a complete architecture and infrastructure, such as BA, for a building-occupant interface. The architecture must be considered in early project planning to ensure that hardware such as submeters are in the purview of the design and construction team. Also, owners must account for the human and computing resources necessary to create as-built asset models and ongoing operational models. A final construction or asset model is a sufficient tool for representing some end uses, but an operational model is needed for the high energy-saving potential systems that are occupant or automatically controllable.

The BA dashboard moved the RSF operational practice from presenting non-contextualized building performance to a first-order data analysis with boundary conditions (presenting end uses with predicted ranges of acceptable energy use). Future work will include efforts to:

- Automate "the art" of building expert-based generation of reduced-order model variable selection
- Include occupant comfort boundaries as well as energy boundaries
- Develop and implement fully automated models with agent reasoning, and proactive comfort maintenance goals and energy achievement goals
- Integrate BA architecture definition into the pre-design phase energy assurance process

The ultimate goal of the BA interface is a tool that extends the concept of integrated design to integrated operations, accounting for many operational factors such as energy price structures and varying occupant preferences.

ACKNOWLEDGMENT

The authors wish to thank NREL sponsors including Ben Kroposki, Stuart Macmillan, and Mark Ruth. In addition, we would like to thank Chad Lobato for his contributions to the initial energy performance assurance activities for the RSF, Gregor Henze of the University of Colorado for his continued support of automated model selection, and the entire Building Agent Team including Kyle Benne, Larry Brackney, Katherine Fleming, John Krueger, Nick Muerdter, Anya Petersen, and Alex Swindler. We would also like to extend a general acknowledgment to the many NREL campus integrated design and construction team members for their contributions to the NREL building procurement and performance assurance process and BA architecture.

REFERENCES

DOE. (2010). 2010 Buildings Energy Data Book. Washington, DC: U.S. Department of Energy, Office of Energy Efficiency and Renewable Energy

Leach, M.; Lobato, C.; Hirsch, A.; Pless, S.; Torcellini, P. (2010). Technical Support Document: Strategies for 50% Energy Savings in Large Office Buildings. 163 pp.; NREL Report No. TP-550-49213. http://www.nrel.gov/docs/fy10osti/49213.pdf

Lobato, C.; Pless, S.; Sheppy, M.; Torcellini, P. (2011). Reducing Plug and Process Loads for a Large Scale, Low Energy Office Building: NREL's Research Support Facility; Presented at ASHRAE Winter Conference, 2011. Preprint. 13 pp.; NREL Report No. CP-5500-49002. http://www.nrel.gov/docs/fy11osti/49002.pdf

Pless, S.; Torcellini, P.A.; Shelton, D.; 2011. Using an Energy Performance Based Design-Build Process to Procure a Large Scale Low-Energy Building. Presented at ASHRAE Winter Conference, 2011. Preprint. 11 pp.; NREL Report No. CP-5500-51323. http://www.nrel.gov/docs/fy11osti/51323.pdf

Schott, M.; Long, N.; Scheib, J.; 2012. "Progress on Enabling an Interactive Conversation between Commercial Building Occupants and Their Building to Improve Comfort and Energy Efficiency" In Proceedings of the 2008 ACEEE Summer Study on Energy Efficiency in Buildings. Monterey, Calif. http://www.aceee.org/files/proceedings/2012/data/papers/0193-000234.pdf

Sheppy, M.; Beach, A.; Pless, S. (2013). Metering Best Practices Applied in the National Renewable Energy Laboratory's Research Support Facility: A Primer to the 2011 Measured and Modeled Energy Consumption Datasets. 16 pp.; NREL Report No. TP-5500-57785. http://www.nrel.gov/docs/fy13osti/57785.pdf