

Cloud-Based Model Calibration Using OpenStudio

Elaine Hale¹, Lars Lisell¹, David Goldwasser¹, Daniel Macumber¹, Jesse Dean¹,
Ian Metzger¹, Andrew Parker¹, Nicholas Long¹, Brian Ball¹, Marjorie Schott¹,
Evan Weaver¹, Larry Brackney¹

¹National Renewable Energy Laboratory, Golden, Colorado

Abstract

OpenStudio is a free, open source Software Development Kit (SDK) and application suite for performing building energy modeling and analysis. The OpenStudio Parametric Analysis Tool (PAT) was extended to allow cloud-based simulation of multiple OpenStudio models parametrically related to a baseline model. This paper describes the new cloud-based simulation functionality and presents a model calibration case study. Calibration is initiated by entering actual monthly utility bill data into the baseline model. Multiple parameters are then varied over multiple iterations to reduce the difference between actual energy consumption and model simulation results, as calculated and visualized by billing period and by fuel type. Simulations are performed in parallel using the Amazon Elastic Compute Cloud service. This paper highlights model parameterizations (measures) used for calibration, but the same multinodal computing architecture is available for other purposes; for example, recommending combinations of retrofit energy saving measures using the calibrated model as the new baseline.

1 Introduction

To prevent climate change and to save money, institutionally and as individuals, we have started expecting our built environment to continually become more energy efficient. This amounts to periodically auditing the energy use of individual buildings via a five-part process: (1) make a basic determination of how much energy the building as a whole, and perhaps individual pieces of equipment, use; (2) investigate the main drivers of that energy use; (3) explore ways the building and its occupants might use less energy; (4) choose and implement changes in equipment and/or operation; and (5) evaluate the efficacy of our actions and repeat this process as needed.

This paper discusses the design and use of software tools that support the first two steps of this process; i.e., quantifying the energy used by a building, and then finalizing a simulation model that can be used to trace energy use to specific end uses and pieces of equipment. The software tools can also support steps 3 and 4: exploring and helping stakeholders choose between various retrofit measures, see Macumber et al. (2014). The overall process can be undertaken at various levels of detail. The focus here is on an ASHRAE level II energy audit, where the baseline is represented by a simulation model calibrated to monthly utility data, and we seek to identify basic retrofit measures (ASHRAE 2004).

The calibration of building energy models to utility and other types of performance data has been an active area of research for over 20 years (Carroll and Hitchcock 1993; Clarke et al. 1993; Reddy 2006). The area remains active because it has proven difficult to build tools and techniques that do not rely on an unaffordable level of expertise to identify model parameters that should be calibrated, and in the end verify that the calibrated model end use breakdowns are likely to be reasonable compared to the actual building. This issue can be decisive-

ly overcome with sufficient hourly or subhourly submetered (by end use) data, and significant computation time; but that solution remains impractical for all but special situations (New et al. 2012). If an auditor has limited performance data, some structure, leading to computational savings, can be imposed on the workflow by breaking the calibration down into steps aligned with weather-independent and -dependent parts of the model (Westphal et al. 2005; Yoon et al. 2003). Sensitivity analysis also plays a large role in trying to keep the number of calibrated parameters reasonable compared to the limited performance data to be matched by the model results. Even so, the temptation to over-fit and over-simulate is strong. For instance, Reddy, Maor, and Panjapornpon (2007) fit models using 20-24 parameters where Sun and Reddy (2006) would recommend 5-6 parameters, overcoming this issue somewhat by recommending retrofits based on the 20 best-calibrated models rather than on a single “solution.”

Bayesian calibration addresses the overfitting problem and provides outputs in support of uncertainty quantification of retrofit energy conservation measures (ECMs) (Heo et al. 2012; Heo et al. 2014). The mathematical approach can be applied to fully detailed (EnergyPlus) models, but is more often used with normative models to keep the computational burden manageable. The main insight is that for the cost of requiring auditors to document uncertainty information for each input, Bayesian statistical methods can be used to transform the auditor’s rough bounds (prior distributions) on parameters into posterior distributions that reflect the likely parameter value given the available evidence. The fact that the prior distributions are modified only as the evidence requires is one form of regularization; i.e., trusting the initial model as much as possible in the face of an underdetermined (more variables than measurements) problem. Even so, the method “...still relies heavily on experts’ judgements especially in the choice of calibration parameters, quantification of their prior distribution, and quantification of uncertainties in other parameters. This will remain so until a repository for standard estimates of parameter uncertainty for [a] variety of building cases becomes available” (Heo et al. 2012).

Without (yet) the ability to gather uncertainty information from auditors, and the limited data available from a level II audit, we conceptualize model calibration as follows:

1. Gather data (detailed geometry and equipment counts, rough schedule information) in the field.
2. Construct an energy model and run the simulation.
3. Note areas of disagreement between the simulation and energy consumption data.
4. Look for gross errors in the initial data, how the data were represented in the model, or areas where there is higher uncertainty in the initial data.
5. Make adjustments and iteratively repeat the process until a satisfactory outcome is achieved (a model whose monthly energy performance matches the utility data with an acceptable amount of error as outlined in ASHRAE Guideline 14).

Such a workflow does not fully validate the accuracy of the final model. However, this process is supported by a standardized method for constructing the initial model (simuwatt Energy Auditor), and by standardized workflows and visualizations in the OpenStudio application suite; discourages overfitting; and can be extended when more data (performance or uncertainty) become available. Thus, the workflow represents an important step toward making simulation-based audits more affordable and common.

Section 2 elaborates on our approach to calibrating building energy models to monthly utility data; Section 3 describes the software tools built to support users in this workflow; Section 4 dives deeply into a case study associated with a demonstration performed for the US Department of Defense’s (DoD) Environmental Security Technology Certification Program (ESTCP); and Section 5 summarizes next steps.

2 Hand Calibration to Monthly Utility Data

Comprehensive ASHRAE levels II and III energy audits require a calibrated energy model to accurately predict the energy savings of each ECM (ASHRAE 2004). Energy model calibration requirements are outlined in ASHRAE Guideline 14, and are applied to electricity use (kWh), electricity demand (kW), and all other fuels used in the facility, as data are available (ASHRAE 2002). The baseline energy model is typically calibrated to monthly utility bills using actual meteorological year (AMY) weather data. ASHRAE Guideline 14 uses two statistical indices to characterize the accuracy of the calibrated model: the mean bias error (MBE) and the coefficient of variation of the root mean squared error CV(RMSE). The MBE indicates how well the model predicts energy consumption as compared to the measured data. Positive values indicate that the model overpredicts actual values; negative values indicate that the model underpredicts actual values. RMSE values indicate the overall uncertainty in the prediction of whole-building energy use. Using monthly data the MBE must be $\pm 5\%$ and the RMSE must be $\leq 15\%$. Energy auditors typically calibrate the baseline model to slightly underpredict energy use so the energy savings predictions are less likely to overestimate energy savings. The amount of data collected during an energy audit depends on the level of the energy audit, the sophistication of the building automation system (BAS), and the types of onsite metering systems (for electricity and natural gas). Assuming an ASHRAE level II energy audit was conducted at a facility with a standard BAS, the following input parameters are typically collected:

- Building geometry and orientation
- Building operational schedules
- Lighting fixture and electric equipment types and counts by space
- Occupant density by space or space type
- Window and door construction and window to wall ratios
- HVAC system characteristics, capacities, zoning, set point temperatures, operational schedules, and control sequences
- Building construction characteristics (if up-to-date drawings are available).
- Utility data for each fuel type
- Process loads

Typically a number of these input parameters are only partially characterized within a certain range of possible values, especially lighting schedules, plug load schedules, specific building construction characteristics (in some cases), infiltration schedules, and various HVAC operation and maintenance issues. Thus, the precision of the entire process is limited by the precision of the input data, and some of these inputs invariably need to be adjusted in a calibration process. Complicating matters, even though hundreds, or thousands (depending on how schedules are parameterized) of parameters are required to fully instantiate a building energy model, it is common for monthly utility bill data to be the only energy use data available. For a single year, this typically amounts to 12, 24, or 36 data points; i.e., 12 readings each for electrical energy, electricity demand, and heating fuel. Mathematically, this yields an incredibly underdetermined problem (many more variables than measurements).

Why not Optimization?

Although calibration (also called inverse modeling or data fitting) is often formulated as an optimization problem, that approach runs into some difficulty because one could easily specify many more parameters to be tuned than there are data points to match. This highly underdetermined problem, if formulated and solved blindly, could easily lead to a situation in

which the utility data are matched within the ASHRAE Guideline 14 tolerance range (“good” R^2 statistic), but the model was tuned such that it does not match reality. If, for example, the good fit was obtained by inferring a poor coefficient of performance (COP), when infiltration was high or an economizer was stuck, an HVAC retrofit ECM might be recommended but will likely provide lackluster results. Also, schedules tend to contribute a good deal of uncertainty, but 1-3 data points a month (electricity use, electricity demand, and heating fuel use) are too few to capture any but the bluntest of schedule errors.

Such a situation can be mitigated in the optimization framework by applying a Bayesian, or other regularization, methodology that biases the resulting model toward initial expectations (prior distributions). However, even doing that does not come for free, as data on the confidence we have in each piece of our model (prior distributions in the Bayesian framework) must be gathered, validated, and subsequently used. Again, this is not an insurmountable barrier, but is highly data intensive. A longer term goal is to produce a semi-automated, robust process for model calibration, but for now we focus on enabling a simple hand-calibration workflow aimed at discovering and fixing gross errors in input data, guided by well-worn error metrics and visualizations that break down the model versus actual data for each billing period and billing metric (energy or maximum demand by fuel type).

Hand Calibration Experience and Methodology

Practical experience performing these sorts of calibrations has taught us to diagnose discrepancies in terms of gross errors; i.e., finding and correcting default or low-confidence inputs that are wrong and significantly affect the model results. These items are very project dependent. In support of this work, we developed a list of the items that are commonly adjusted in the hand-calibration process (see Table 1).

Table 1. Common Adjustments Made During Hand-Calibration to Monthly Utility Data

Measure Name	Description and Comments	Importance
Adjust Infiltration	Adjust per-area infiltration. Adjust infiltration or outside air, not both. Use default unless building seems extra-leaky.	Primary
Adjust Outside Air	Adjust the mechanical ventilation provided to a space. Usually listed correctly on newer buildings’ drawings, but often uncertain in older buildings, or affected by damper faults.	Primary
Occupancy Sensor Savings	Adjust the savings associated with occupancy sensors in a certain space. Inspection of actual operation may justify a deviation from the ASHRAE default credit of 10%.	Secondary
Heating and Cooling Temperature Set Points	Adjust the thermal zone air temperature set point for the heating or cooling system. As an easily adjustable operational variable, this may be uncertain.	Secondary
Supply Air Temperature	Adjust the supply air temperature, possibly in support of a seasonal reset strategy. Typically these are constant values, but in some cases the BAS shows a range.	Secondary
Chilled and Hot Water Temperature Set Points	Adjust the temperature at which chilled water and hot water are supplied (typically to one or more air systems). These are typically constant, but certain conditions (inaccessible controls, BAS shows a range) may introduce uncertainty.	Secondary
Economizer Set Points	Adjust the economizer set points. May be uncertain if the control system is inaccessible.	Secondary
Fan and Pump	Adjust the fan or pump motor efficiency. This is likely to be	Secondary

Motor Efficiencies	correct if each fan or pump is modeled individually, but can be harder to get right if a whole-building average is used.	
Ground Temperatures	Set the temperature of the ground in immediate contact with the building to the average indoor air temperature minus 3.6°F (2.0°C) or use a simulation program to model specific building geometry, soil conditions, and envelope insulation.	Secondary
Adjust Plug, Lighting, or Exterior Lighting Loads or Schedules	Adjust peak loads or schedules across the board, or by equipment type, space type, or individual space. Schedules and process loads can be hard to capture correctly without submetering. The modeling methodology may have a weak point, especially as applied to specific buildings.	Tertiary
Boiler or Chiller Efficiency	Adjust the boiler or chiller efficiency. This is usually well known for new-enough equipment, but it can be hard to get a good number for old equipment.	Tertiary
Hot Water Usage and Supply Temperature	Adjust the service hot water usage and/or supply temperature. In select space types (e.g., gyms), the amount of hot water use and its temperature is significant. The latter may not have been read off the tank.	Tertiary
Insulation R-Value	Adjust insulation R-values in the exterior floor, walls, or roof. Sometimes the thickness of the insulation is known, but the material is not. This detail can be important in very cold climates.	Tertiary
Adjust Window Properties	Adjust the U-value and solar heat gain coefficient for the windows. It can be hard to get good numbers for these properties for very old buildings.	Tertiary
Adjust Operating Hours	Move the opening and closing times of the building (and related schedules) for special times of the year not caught in initial discussions with the site staff.	Tertiary

3 Hand Calibration with OpenStudio

A software-guided workflow enabling hand calibration of energy models has been available since version 1.1 of the OpenStudio SDK and application suite. It is made more efficient and reliable through the use of OpenStudio measures (self-contained Ruby scripts that transform models in specific ways) corresponding to items described in the previous table. These “calibration measures” are made available via the Building Component Library (BCL), an on-line repository of modeling content that is tightly integrated with OpenStudio. To enable this workflow, users need to specify some utility bill data and an AMY weather file in their baseline models. This triggers the SDK to compute calibration metrics and generate reports, graphical and tabular, to show the user how well their model matches the given utility data.

The initial model may not match the utility data within an acceptable amount of error, at which point the user should review the data and the calibration outputs to hypothesize where any large errors (observable in utility bill data) may reside. The user can then try to correct those errors, either in the OpenStudio Application, or, if there are several possible sources of interacting errors, in PAT, which provides side-by-side comparisons of multiple scenarios. We will now outline the workflows in both applications in greater detail.

OpenStudio Application

The user first creates an OpenStudio model of the building using the OpenStudio Application (combined with the Sketch-Up Plugin) or another software tool (e.g., simuwatt Energy Auditor). This initial model reflects the data collected during an audit. In addition to the model, calibration requires utility bill data and an AMY weather file covering the time period of the utility data. The workflow follows:

1. Add utility data for various fuel sources to cover the entire calendar year. This will likely result in 13 utility periods starting in December of the year previous to the AMY file, and ending in the January following the AMY file.
2. Load the AMY weather file covering the time periods in the utility data.
3. Change the calendar year of the model to match the AMY file.¹

After the baseline model is constructed, it can be simulated with EnergyPlus. The utility bill data and an AMY weather file automatically trigger a calibration report that graphically shows the actual and utility bill energy use and demand (if available) by fuel type, along with the ASHRAE Guideline 14 metrics (see Section 2). The user then examines these data to see (1) if the error between the model and the utility data is sufficiently small and (2) if not, where the major errors may lie based on engineering judgement applied to the actual and modeled energy use and demand patterns. For example: Is gas use low in the winter? Is electricity use high in the summer, or all year long? The answers to such questions are clues to deficient aspects of the model.

Parametric Analysis Tool

From the OpenStudio Application calibration report, the user may have several ideas for where gross errors lie. If this list is more than a few items long, or if the EnergyPlus model takes a long time to run, using PAT to explore the calibration possibilities may be useful. This workflow follows:

1. Start a new PAT project and specify the baseline model.
2. Find (on the BCL) or write custom calibration measures that can be used to model the suspected errors.
3. Apply the measures to the project via PAT's drag-and-drop interface on Tab 1 and specify associated input values (e.g., COP or cooling set point).
4. Use Tab 2 to specify which combinations of parameter changes to consider.
5. Run models locally (for small parameter sets or while debugging) or on the cloud on Tab 3.
6. Examine results (PAT Tab 4) compared to the baseline model using the same annual calibration metrics calculated in the Application.
7. Select a "calibrated" model for further retrofit analysis.

This workflow will be demonstrated more concretely in Section 4.

Running on the Cloud

Since OpenStudio version 1.1, design alternatives made with PAT may be run using the Amazon Web Services (AWS) Elastic Compute Cloud (EC2) service. To do this, the user sets up an AWS account and then connects to it through PAT. Once it is configured, the user can

¹ If the simulation is to run a full year, and if the utility bills don't exactly cover the AMY year, OpenStudio will calibrate the utility periods that span January 1st using results from December and January of the AMY year. If this approach is not acceptable, setting the run period to January 2nd through December 30th on the simulation settings tab will exclude the first and last utility periods from the calibration.

“Start the Cloud” on Tab 3. This initializes server and worker nodes with preconfigured images that include OpenStudio, EnergyPlus, and other required software. Once the nodes are ready, pressing the run button transmits the baseline model, measures, and parameters to the server, which distributes parallel tasks to worker nodes. The server aggregates high-level results that are returned to the PAT client application. The user may also download complete simulation results for each calibration alternative.

4 Case Study

This case study was generated as part of an energy audit performed on a 7,560-m² (81,400-ft²) office building at Tyndall Air Force Base in Panama City, FL. The building is all electric and uses a chilled water variable air volume system with electric reheat for space conditioning. The occupancy is primarily office space with small areas that include a courtroom, restrooms, conference rooms, a health clinic, and mechanical space. The system has not been commissioned recently, and occupants complained of hot and cold spots. The building was constructed in 1988, and since has had many space use changes, as well as several wall demolition/construction projects that did not include HVAC redesign. The model was calibrated using AMY data appropriate for the site.

The 2012 utility data show a building energy use intensity of 298 MJ/m² (26.2 kBtu/ft²). In this cooling-dominated climate, such low energy use indicates a reasonably efficient HVAC system and low to moderate plug loads. After the model was built, zoned according to the HVAC drawings, and had the loads added to the space, it showed much higher electricity use than the actual data, with a -66.89% electricity normalized mean bias error (NMBE), see Figure 1.

The model was calibrated iteratively over several rounds. The first iteration focused on ensuring ground temperature assumptions were reasonable for the climate, and checking modeled versus known pump and fan motor characteristics and efficiencies. The latter were in good agreement. The ground temperatures, however, required more attention.

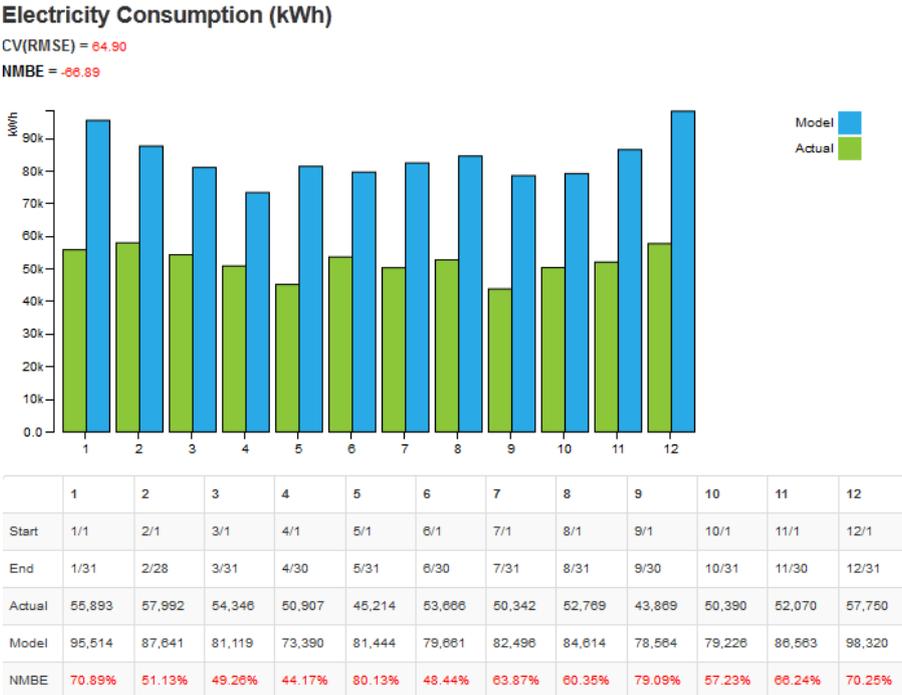


Figure 1. Calibration report for electricity in the uncalibrated model

Electricity Consumption

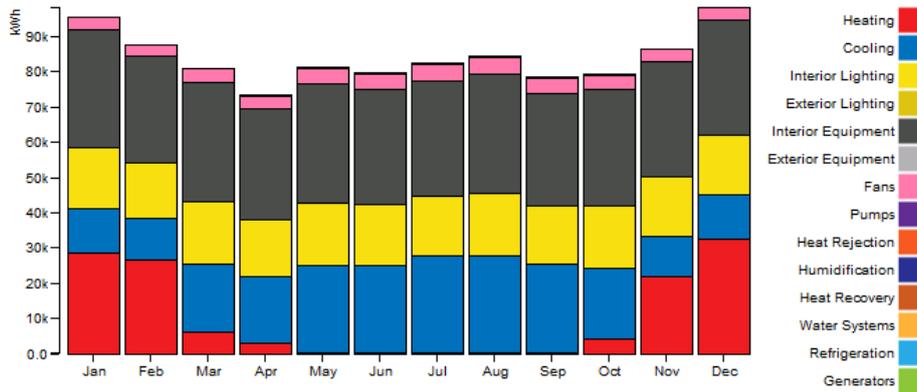


Figure 2. End use breakdown for the uncalibrated model

The initial model used the EnergyPlus default assumption of constant 18.0°C (64.4°F) ground temperatures. According to the rule of thumb that the ground temperature under commercial buildings is approximately equal to the average indoor space temperature minus 2.0°C (3.6°F), this would correspond to an average indoor temperature of 20.0°C (68.0°F), which seemed too cold for a Florida office building. To better approximate the ground temperature, we wrote a script that retrieves the indoor air temperature time series from each first floor thermal zone, and uses those data to calculate the floor area-weighted average indoor temperature over the entire model year. For the initial model, that temperature was 20.9°C (69.6°F), so an OpenStudio measure for setting the ground temperatures to a constant value for the year was used in the subsequent round to set the ground temperatures to 18.9°C (66.0°F).

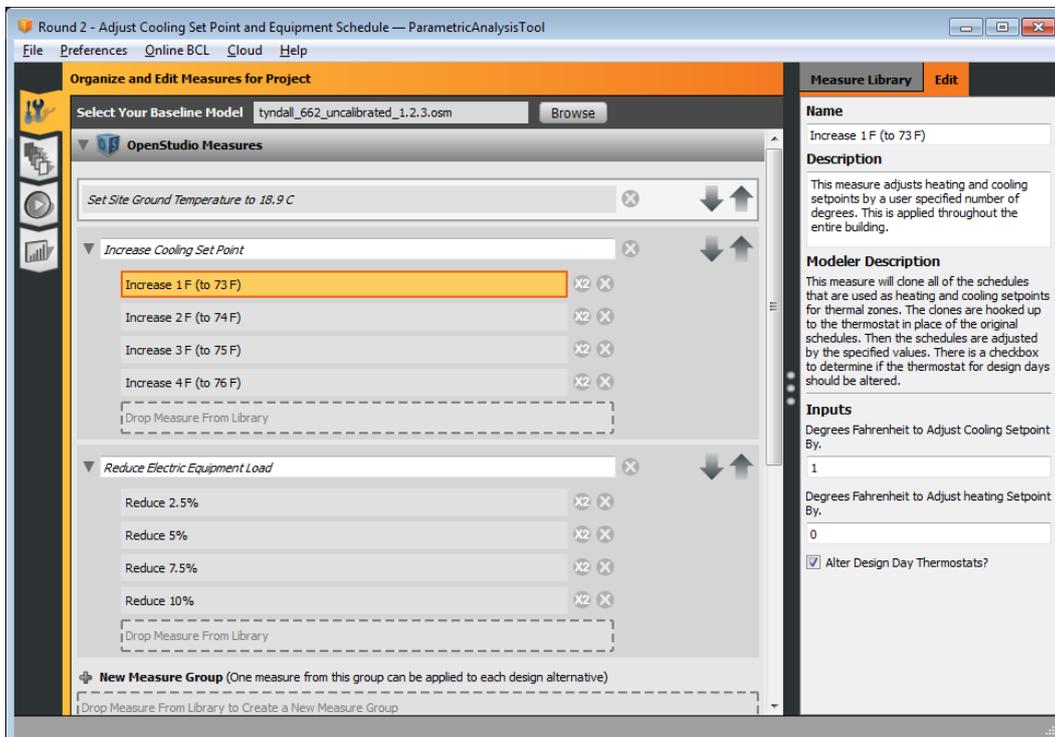


Figure 3. Calibration measures selected and configured

We then focused on why the model energy use was much higher than the actual energy use. From the end use breakdown of the uncalibrated model (Figure 2), and based on occupant comments that parts of the building often feel too warm, we reduced plug load consumption and increased the cooling set point, which in the initial model was 22.2°C (72.0°F).

While such changes can be made manually in the OpenStudio Application, PAT can be used to systematically evaluate a range of parameters. Also, the ability to run simulations on cloud resources expedites trying out different calibration parameters. For example, this model takes about 15 minutes of EnergyPlus simulation time. A personal computer can typically run only 2-8 simulations simultaneously. By using the PAT cloud feature with one server and one or more worker nodes, the number of simultaneous simulations can be vastly increased at a small cost. At the time of this writing, a user could run 64 simulations in parallel for about \$12/hr, giving a throughput of 256 simulations per hour (compared to 8-64 simulations per hour for a typical computer) for a fraction of the cost of an energy modeler’s time. For this first major round of calibration, we used PAT with cloud computing resources to run the model with several combinations of cooling set point changes and plug load reductions.

In particular, the project-specific ground temperature measure was selected and its only argument, the constant ground temperature to be applied for the whole year, was set to 18.9°C (66.0°F). Increasing the cooling set point and reducing the plug loads could be modeled in several ways. In the interest of time, we repurposed two measures already available from the BCL. The cooling set point was adjusted using the “Adjust Thermostat Set Points by Degrees” measure, setting the adjustment to the heating set point to zero, and looking at four alternatives for the cooling set point (increases of 1, 2, 3, or 4°F, that is, cooling set points of 73-76°F). Plug loads were reduced across the board with the “Reduce Electric Equipment Loads by Percentage” measure. Because the plug loads were individually audited, from a modeling perspective it would have been better to lower the electric equipment *schedule* by a

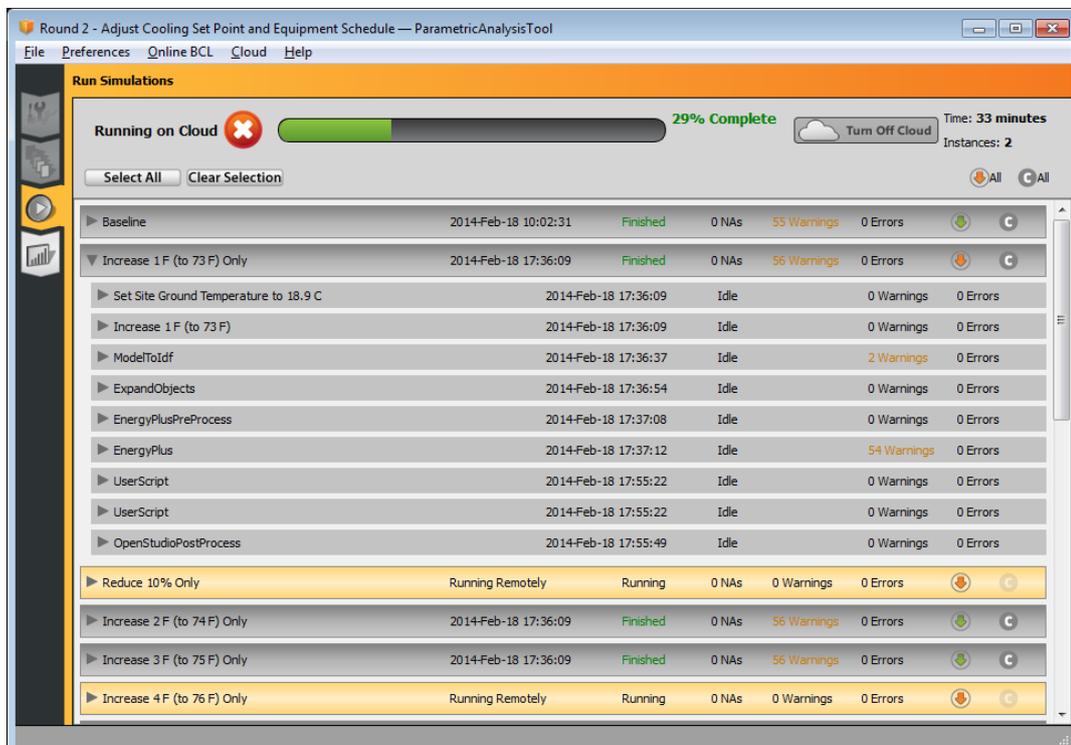


Figure 4. Parametric analysis tool run tab while simulations are running on the cloud

percentage, rather than the load level, but this route was expedient. The final configuration is shown in Figure 3.

All combinations of baseline values and possible changes to the cooling set point and the electric equipment load were used to create 25 unique design alternatives. This number of simulations would take about an hour to run on our typical 8-processor laptops, so the baseline model was first run locally, and then the remaining 24 points were run on the cloud (Figure 4). With a single 16-processor worker node, just two sets of simulations were required, thus saving about 15 minutes in clock time compared to running all simulations locally. Sending those simulations to the cloud also freed the local computer for other tasks. While running on the cloud, PAT continually monitors progress, updates design alternatives with summary results as they become available, and downloads full simulation results for points for which they have been requested. The AWS console may also be used to get the server’s public DNS by navigating to the running instances on the EC2 dashboard, selecting the OpenStudio-Server instance, and then copying it from the header of the bottom pane. Pointing your browser to that URL takes you to the OpenStudio Cloud Management console, which enables monitoring and cluster control. The console is shown midrun in Figure 5.

From this set of runs, we learned that even with the cooling temperature set point at 76°F and 10% reduced electricity loads, the model energy use was still too high compared to the actual utility data. Assuming that the cooling set points are reasonably close, a more thorough examination of the model suggested that improvements to the lighting and electric equipment schedules were warranted. For example, the initial electric equipment schedule is appropriate for computer-heavy buildings, which often have high values at night (0.3) and heavy use during the day; however, the audit data for this building indicated that electric equipment is not dominated by computers, but also contains vending machines, coffee makers, TVs, monitors, printers, microwaves, refrigerators, projectors, and water coolers, which have different schedule characteristics.

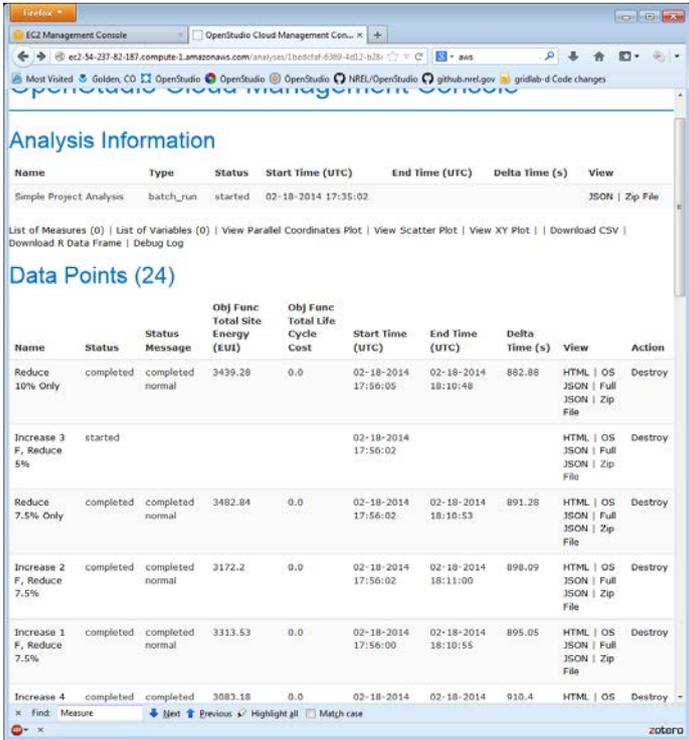


Figure 5. OpenStudio Cloud Management Console during run initiated from PAT

After reconfirming that the ground temperature should be warm (and raising the ground temperature to 19.9°C (67.8°F) to reflect the higher cooling temperature set point), we made the following changes to the model in the OpenStudio application:

- Adjusted the lighting schedule to reduce the nighttime and weekend fraction to 0.1, and the highest fractional level to 0.6 based on the heuristic that office building loads are typically proportional to occupancy, and occupancy is usually 60-70% (GSA 2008).
- Similarly adjusted the electric equipment schedule, because the plug loads were diverse and individual pieces of equipment were modeled with generic (typically high-power) prototypes.
- In response to undersizing by EnergyPlus’s autosizing routines (compared to the actual size of the fan motors), we hard-sized the two constant volume fans based on the guideline of 0.00508 m³/s·m² (1.00 cfm/ft²).

PAT was also used to sweep the fan static pressure between 500 and 870 Pa (2.0 and 3.5 in. H₂O). This resulted in motor sizes produced by EnergyPlus that were in much closer agreement with actual audited nameplate data.

All of these changes, along with an across-the-board reduction in electric loads by 40%, produced calibration results that agree well for the shoulder seasons and summer. The corresponding end use plots show a good deal of cooling and reheat November through March that could easily be avoided with a wintertime reset strategy. As the chiller control details were not available during the audit, and could therefore have included a reset strategy, we decided to model this by setting the supply air temperature to 15.6°C (60.0°F) November 1 – March 15. This change brought about the calibrated model results shown in Figure 8 and Figure 9.

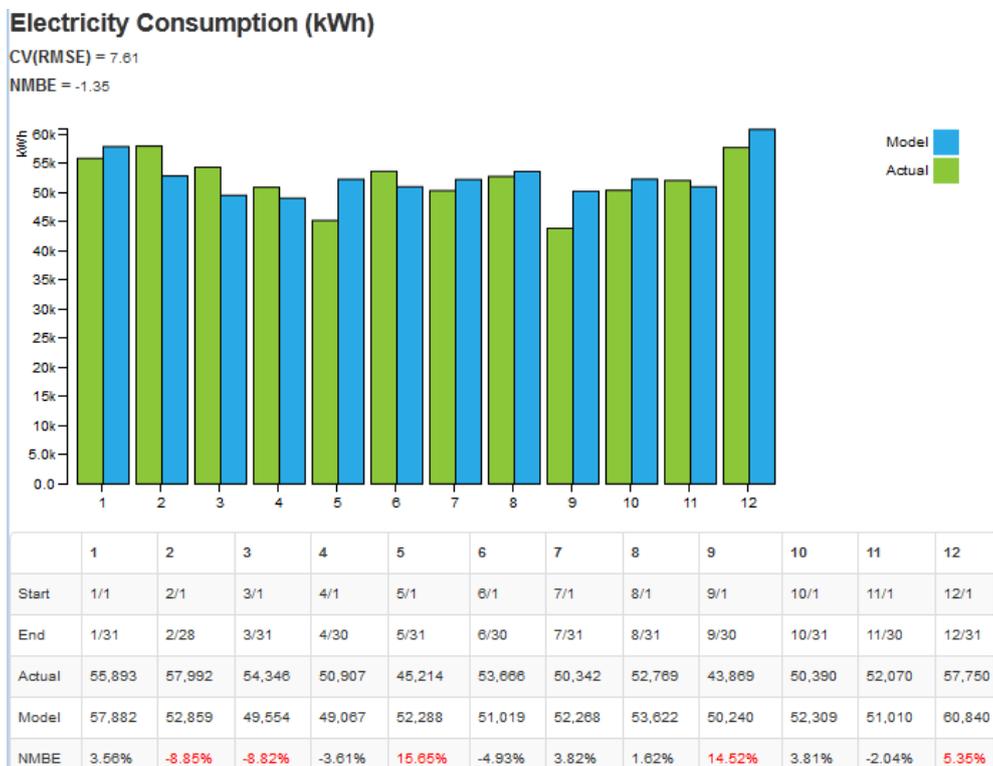


Figure 6. Final calibration results

Electricity Consumption

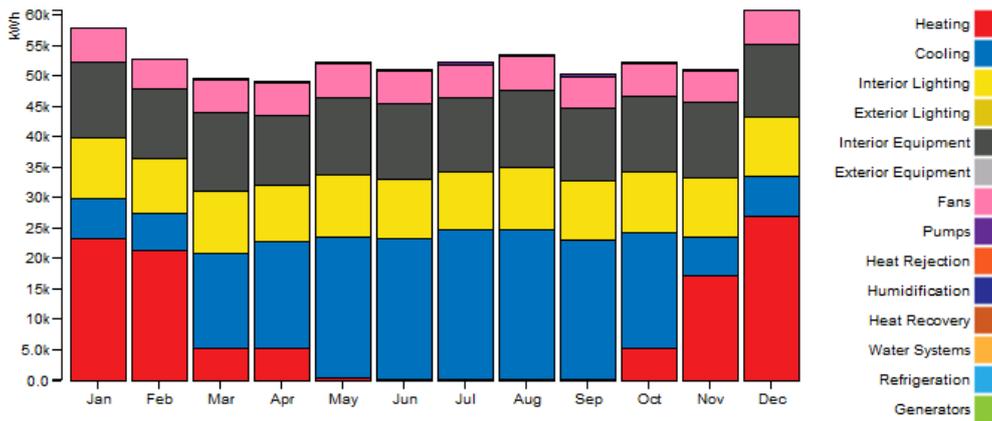


Figure 7. End use breakdown for the final, calibrated model

Note that the May and September deviations remain significant. The utility data on Tyndall Air Force base are gathered internally, and the accuracy of the bins in which we received the data (which line up exactly with the calendar months) may not be wholly accurate. Also, with just 12 electricity use data points, and minimal communication with the building manager, determining if the right calibration measure adjustments have been made is impossible. Known values, rules of thumb, relative levels confidence, and the shapes of the calibration deviations informed our measure and parameter choices. Table 2 shows a summary of our changes and the reasons behind them. More confident and finer adjustments (getting into the overall shape of schedules, for instance) would require additional utility data and more detailed audit data, such as non-generic plug load descriptors, and more diverse schedule sets.

Table 2. Calibration Summary

Change	Rationale
Raise site ground temperature from default of 18.0°C (64.4°F) to 19.9°C (67.82°F).	Apply rule of thumb that ground temperature is 2°C (3.6°F) lower than average indoor air temperature.
Increase cooling set point to 24.4°C (76°F)	Occupants report hot spots and modeled energy use was significantly higher than actual.
Adjust lighting and electric equipment schedules.	Simplified the schedule shapes, and brought the overall values in line with rules of thumb based on occupancy, nighttime lighting levels, and plug load types.
Reduce electric equipment loads by 40%	The diversity of plug loads all applied with the same schedule, and the disproportionate share of plug loads as compared to lighting loads.
Hard size fans to 1.00 cfm/ft ² (0.00508 m ³ /s·m ²)	Autosized fan motors were undersized compared to actual, known values. 1 cfm/ft ² is a rule of thumb.
Increase the fan static pressure to 3.5 in. H ₂ O (870 Pa)	The fans motors were still too small after hardsizing to 1 cfm/ft ² .
Reset the supply air temperature to 15.55°C (60°F) (from 12.77°C [55°F]) in winter	After the shoulder and summer season energy use was brought in line, the winter energy use was still too high. The end use breakdown showed considerable cooling plus reheat, and the controls on the chiller were not fully audited, leaving open the possibility of winter reset.

5 Conclusions

The process of conducting a building energy audit, including model-building, calibration, and making retrofit recommendations, is fairly well known and documented. However, this process is still highly manual and idiosyncratic. This paper describes a tool chain and corresponding workflow for hand-calibrating simulation models to monthly utility data that makes the process faster, less error-prone, more replicable, and more self-documenting.

The process was successfully used to bring a model of an office building, produced with the simuwatt Energy Auditor tool, within the bounds specified by ASHRAE Guideline 14. This case study helps highlight what is possible, what is not possible, and which data are especially important in trying to calibrate a building model, which has hundreds to thousands of individual input parameters, using monthly energy data and an AMY weather file.

Of utmost importance are the actual utility data. For the case study model, it is unclear whether the billing period dates were properly assigned for each utility data bin, especially given the large deviations in May and September. Other case studies undertaken by the larger project also had problems with their utility data including:

- Missing data for natural gas
- Utility meters that covered multiple buildings
- Unmetered chilled and hot water from central plants
- Data that seemed abnormally high/low
- Peaks in months that could not be explained by environmental conditions/activities

Monthly utility data do not provide the granularity necessary to distinguish between mismodeled lights and mismodeled computers, for instance. Schedule discrepancies are invisible without daily, or better yet, hourly data. For these and other reasons, we are heartened by the US Federal Energy Management Program and DoD efforts to collect subhourly data on many buildings. Furthermore, the hand-calibration process is quite labor intensive. Numerical optimization methods promise to ease this burden, but without larger quantities of utility data (daily or hourly by fuel type), and/or the ability for auditors to easily enter uncertainty descriptions for each input, the promise will be difficult to realize.

Current work includes supplementing this manual process with more automated, optimization-based calibration processes that are used in conjunction with data visualization and quality control/quality assurance procedures to help modelers avoid common calibration pitfalls, such as adjusting more input parameters than the amount of data available would recommend or simultaneously adjusting parameters whose impacts on the model are indistinguishable from one another.

6 Acknowledgements

The authors wish to thank the DoD ESTCP for supporting this work as well as all the site personnel who assisted with the site audit that provided the basis for this case study. We appreciate continued support from the U.S. Department of Energy's Building Technologies Office, which has produced the underlying OpenStudio platform and BCL. We are also grateful to our eSim conference peer reviewers for their constructive comments on the abstract and an early draft.

7 References

- ASHRAE. 2002. "ASHRAE Guideline 14-2002: Measurement of Energy and Demand Savings". Guideline. Atlanta, Georgia: American Society of Heating, Refrigerating and Air-Conditioning Engineers, Inc.
- . 2004. *Procedures for Commercial Building Energy Audits*. American Society of Heating, Refrigerating and Air-Conditioning Engineers, Inc.
- Carroll, W.L., and R.J. Hitchcock. 1993. "Tuning Simulated Building Descriptions to Match Actual Utility Data: Methods and Implementation." *ASHRAE Transactions* 99 (2): 928–34.
- Clarke, JA, PA Strachan, and C. Pernot. 1993. "An Approach to the Calibration of Building Energy Simulation Models." *TRANSACTIONS-AMERICAN SOCIETY OF HEATING REFRIGERATING AND AIR CONDITIONING ENGINEERS* 99: 917–917.
- GSA. 2008. "Sustainability Matters". Washington, D.C.: U.S. General Services Administration. http://www.gsa.gov/graphics/pbs/Sustainability_Matters_508.pdf.
- Heo, Y., R. Choudhary, and G. A. Augenbroe. 2012. "Calibration of Building Energy Models for Retrofit Analysis under Uncertainty." *Energy and Buildings* 47: 550–60. doi:10.1016/j.enbuild.2011.12.029.
- Heo, Yeonsook, Diane J. Graziano, Leah B. Guzowski, and Ralph T. Muehleisen. 2014. "Evaluation of the Efficacy of Bayesian Calibration: Interim Project Report". Technical Report ANL/DIS-14/1. Chicago, Illinois: Argonne National Laboratory.
- Macumber, Daniel, Andrew Parker, Lars Lisell, Matthew Brown, and Ian Metzger. 2014. "Simuwatt -- A Tablet Based Electronic Auditing Tool." In *Proceedings of eSim 2014*. Ottawa, Ontario, Canada.
- New, Joshua R., Jibonananda Sanyal, Mahabir Bhandari, and Som Shrestha. 2012. "Autotune E+ Building Energy Models." *Proceedings of the 5th National SimBuild of IBPSA-USA*: 1–3.
- Reddy, T. Agami. 2006. "Literature Review on Calibration of Building Energy Simulation Programs: Uses, Problems, Procedures, Uncertainty, and Tools." *ASHRAE Transactions* 112 (1): 226–40.
- Reddy, T. Agami, Itzhak Maor, and Chanin Panjapornpon. 2007. "Calibrating Detailed Building Energy Simulation Programs with Measured Data - Part II: Application, to Three Case Study Office Buildings (RP-1051)." *Hvac&r Research* 13 (2): 243–65. doi:10.1080/10789669.2007.10390953.
- Sun, J., and T.A. Reddy. 2006. "Calibration of Building Energy Simulation Programs Using the Analytic Optimization Approach (RP-1051)." *HVAC&R Research* 12 (1): 177–96.
- Westphal, F.S., R. Lamberts, and F.S. Catarina-Brazil. 2005. "Building Simulation Calibration Using Sensitivity Analysis." In *Building Simulation*, 9:1331–38.
- Yoon, Jongho, E. J. Lee, and D. E. Claridge. 2003. "Calibration Procedure for Energy Performance Simulation of a Commercial Building." *Journal of Solar Energy Engineering* 125 (3): 251–57. doi:10.1115/1.1564076.