

Uncertainty Analysis of a Heavily Instrumented Building at Different Scales of Simulation

George OSTROUCHOV¹, Joshua NEW², Jibonananda SANYAL^{3*}, Pragneshkumar PATEL⁴

Oak Ridge National Laboratory

¹Computer Science and Mathematics Division

^{2,3}Building Technology Research & Integration Center

⁴Joint Institute for Computational Sciences

One Bethel Valley Road, Oak Ridge, Tennessee, United States

Phone: 865-574-3137, Email: ostrouchovg@ornl.gov

Phone: 865-241-8783, Email: newj@ornl.gov

Phone: 865-241-5388, Email: sanyalj@ornl.gov

Phone: 865-576-0627, Email: pragnesh@utk.edu

ABSTRACT

Simulation plays a big role in understanding the behavior of building envelopes. With the increasing availability of computational resources, it is feasible to conduct parametric simulations for applications such as software model calibration, building control optimization, or fault detection and diagnostics. In this paper, we present an uncertainty exploration of two types of buildings: a) of a building envelope's thermal conductivity properties for a heavily instrumented residential building involving more than 200 sensors, and b) a sensitivity analysis of a stand-alone retail building from the U.S. Department of Energy's reference model. A total of 156 input parameters were determined to be important by experts which were then varied using a Markov Order process for the residential building generating hundreds of GBs of data for tens of thousands of simulations. For the commercial building, 20 parameters were varied using a fractional factorial design requiring just 1024 simulations generating data in the order of a few hundred megabytes. These represent a wide variety and range of simulations from a few to tens of thousands of simulations in an ensemble.

Depending on the number of simulations in an ensemble, the techniques employed to meaningfully make sense of the information can be very different, and potentially challenging. Additionally, the method of analysis almost always depends on the experimental design. The Markov Order sampling strategy and fractional factorials designs of sampling presented represent two approaches one could employ for large sensitivity analysis of buildings at two different scales of simulations. The paper presents the analysis using descriptive statistics as well as employing multiple analysis of variance techniques for comparison and contrast.

1. INTRODUCTION

There are many use cases that leverage software simulations to quantify changes to building behavior. For an individual homeowner or building portfolio manager, simulation is often used to apply dozens or hundreds of energy conservation measures (ECMs) to a (set of) building(s) and determine the optimal return on investment. This use case is employed by several governmental agencies such as the Office of Weatherization and Intergovernmental Programs (OWIP) for low-income housing, Federal Energy Management Program (FEMP) for federal facilities, and the \$5 billion energy service company (ESCO) market which often profit primarily by sharing in weather-normalized energy savings over time. For legal tradeoffs, such as those in building codes or taxes, simulation is used to allow flexibility in the way a building is constructed and still meet minimum energy efficiency requirements as well as determine the best ways to increase the energy efficiency of building codes when they are updated. For incentives such as those from local utilities, state governments, or federal policies, simulation is often used to determine eligibility; many incentives can be found in tools online such as the Database for State Incentives for Renewable Energy (DSIRE). For national energy reduction, organizations such as the Department of Energy (DOE)

have been tasked by the U.S. President with aggressive goals to reduce energy waste from buildings 50% by 2030 (compared to a 2010 baseline). DOE relies on simulation and other analysis to optimize policy selection, determine and prioritize funding for high impact technologies (HITs) that enable actualized energy savings, and is used extensively to model, optimize performance, and integrate individual building components (envelope, equipment, etc.).

There are many simulation engines and software tools, which leverage these simulation engines that have made it easier than ever to run complex analysis. Some of these engines and tools are highly specialized and seek only to optimize a constrained set of high-fidelity physics to enable optimization and creation of specific types of roofs, HVAC equipment, etc. However, many of these specialized algorithms have been developed within or absorbed by well-known engines and tools in order to maximize their utility for energy-efficiency experts. Within the context of the Department of Energy, a whole-building simulation engine known as DOE-2.1E was first released in 1993 and is now up to DOE-2.3 that can run an annual building energy simulation in seconds and is still widely used by many experts through tools such as eQuest which allows one to run groups of simulations. In 2001, DOE was beta-testing a whole-building simulation engine known as EnergyPlus that was released later in the same year –and has continued to grow given DOE’s steady investment in its development. The typical runtime of EnergyPlus is on the order of a few minutes to run an annual simulation and this is one of the major issues that have limited its adoption. It is, however, in widespread use today and contains many of the most recent building dynamic simulation algorithms. Many algorithms of varying fidelity exist for modeling certain phenomena within the simulation engine allowing the user to define the tradeoff between more accurate simulations and longer runtime. EnergyPlus consists of ~600,000 lines of Fortran code and has recently been cross-compiled to ~750,000 lines of C for transitioning to a Modelica-based architecture currently being referred to as the Son of EnergyPlus (SOEP).

In 2009, DOE released OpenStudio, which has evolved to become the primary Application Programming Interface (API) for EnergyPlus. It contains many tools, such as the Parametric Analysis Tool (PAT), for applying measures to buildings and producing analytical reports. While the usefulness of longer-running EnergyPlus simulation has been mitigated for deployment on the desktop, OpenStudio now supports the use of cloud computing; this allows anyone in the world to run large-scale simulations and analysis, without the need for servers or supercomputers, at the cost of current cloud computing resources. A large proportion of DOE’s investment in building energy simulation and software tools is captured in EnergyPlus and OpenStudio. There are many other building tools beyond the scope of this paper; for a comparison of engine and tool capabilities, the interested reader should refer to Crawley, 2008.

As computing resources and economically feasible large-scale analysis become more prevalent, there is an increasing need to develop simulation analysis that can scale to the number of cores available. In addition, methods are needed for dealing practical issues such as simulation failures when sampling the combinatorially large parametric space of input parameters. While most modern software tests for valid ranges of inputs for the simulation engine using regression tests or for an individual module with unit tests, virtually no software addresses how a set of parameters in a valid range may cause the software to crash or give unrealistic results when used in different combinations. With respect to EnergyPlus, a standard input file to describe a building has on the order of 3,000 inputs. We’ve had domain experts pick out the most important parameters that affect building energy consumption and meaningfully discretize values for those parameters. The most important group consisted of 156 input parameters and yielded a combination of 5×10^{52} simulations that would be necessary to brute-force every combination. Even if the Department of Energy wished to document input combination that were not viable, the computation time to find the non-linear boundaries of what constituted safe system state would be a tremendous challenge; to say nothing of the challenge for documenting such boundaries. While this may be an extreme case, a more common scenario is when a user runs multiple simulations while tweaking individual parameters, which often results in waiting for failed simulations. For this reason, we have elected to analyze a realistic use case in which naïve parameter selections cause multiple groups of simulations to fail and then compare the analysis when domain experts go back to fix the failed simulations.

2. STATISTICAL TECHNIQUES FOR SENSITIVITY ANALYSIS

2.1 Experimental Design

In statistical design of experiments, input parameters of a multivariate function under study are called *factors* and output parameters are called *responses*. The objective of a designed experiment is to measure the effect sizes of

factors and their interactions on a response through an additive decomposition of the multivariate response function. When many possible factors need to be screened for their influence on a response, two-level fractional factorial designs are popular. In general, for k factors at two levels each, there are 2^k possible factor combinations. Fractional factorial designs allow the use of a small fraction of this large number by carefully controlling the pattern of factor and interaction effect aliasing (see, for example, Mee, 2009). In this study, we have 20 factors and we use a resolution VI design, resulting in 1024 runs. Resolution VI designs alias main effects (single factors) with order five or higher interactions and alias two-factor interactions with order four or higher interactions. We generate the design with the R package FrF2 (Ulrike Grömping, 2014).

2.2 Sensitivity Analysis

The analysis of variance of a resolution VI fractional factorial design can be done with a model that contains all main effects and all two-factor interactions. Each of the 1024 experimentally configured buildings is simulated over a period of one year and we use the 12 monthly totals in our analysis. We use the 12 measurements in a repeated measures multivariate analysis (Hand and Taylor, 1987) with R software (R Development Core Team, 2014). This traditional multivariate approach to repeated measures data is considered more powerful than treating each month as a separate analysis of variance. Within the buildings domain, many of the motivations, history, and state-of-the-art approaches for sensitivity analysis (reducing the number of simulations necessary, optimization methods, etc.) are expertly consolidated in ASHRAE report 1051-RP (Reddy, 2006).

3. EXPERIMENTAL SETUP OF SIMULATIONS

3.1 ZEBRAlliance Homes and Markov Order Design

Measured data was used from the first house from a set of 4 residential buildings known as the ZEBRAlliance (Miller, 2012). This 2800 ft² research home collected 269 channels of 15-minute sensor data during the 2010 calendar year (Biswas, 2012). The reference building used in this work is house number 1 in the Wolf Creek subdivision (WC1), an Oak Ridge National Labs ZEBRAlliance experimental energy efficient home. This home has a plethora of energy-efficient technologies: (1) standing seam metal roof with infrared reflective pigments to boost solar reflectance, (2) ENERGY STAR appliances, (3) triple-pane low emittance Argon-filled windows, (4) compact fluorescent lighting, (5) horizontal ground loop installation that leverages foundation and utility excavations, (6) high-efficiency water-to-air heat pump for space conditioning, (7) high-efficiency water-to-water heat pump for hot water heating, (8) an energy recovery ventilator for transferring heat and moisture between fresh incoming and outgoing air, and (9) structurally insulated panel (SIP) walls filled with expanded polystyrene insulation. All four buildings used emulated occupancy, running the exact same schedule of plug loads, lighting, and hot water draws in the buildings in order to remove the confounding variable of behavior-based variability typically encountered in occupied building comparisons.

After constructing an EnergyPlus model of this building, domain experts identified 156 of the model inputs that they considered the most important variables for sensitivity analysis. A minimum, maximum, and default value were selected for each of these parameters along with a physically-meaningful step size (serving to discretize the search space). From the 269 channels of measured data, 96 outputs were found which corresponded to something EnergyPlus could report. In order to conserve space, the 96 outputs and 156 inputs used have been made available at http://bit.ly/autotune_res_params.

There were three main sampling methods used to vary the input parameters.

1. Brute Force – 14 of the 156 building input parameters were selected for brute-force calculation of all combinations of minimum and maximum sizes. This fine-grained (FG) dataset contains approximately 12,000 EnergyPlus simulation results and is approximately 143GB.
2. Markov Order 1 – all 156 parameters were kept at their default value while iterating through each parameter to set it at its extremes. This has the ability of showing the swing in outputs as a function of a single parameter going from its minimum to its maximum value. This dataset contains 299 simulations and is 3.9GB.
3. Markov Order 2 – all 156 choose 2 pairings were run in all possible min/max configurations (min/min, min/max, max/min, max/max). This overcomes the limitation of Markov Order 1 in that interplay between any pair of variables could be ascertained. This dataset contains approximately 28,000 EnergyPlus simulation results and is approximately 450GB.

3.2 DOE's Reference Stand-Alone Retail and Fractional Factorial Design

DOE has defined a set of 16 commercial reference buildings (Deru, 2011) and 3 vintages for each of the 16 ASHRAE climate zones. These models are used by the building energy modeling community since they directly characterize 60% of the commercial buildings found in the U.S. building stock and are similar to the remaining types. The stand-alone retail building is one of the most prevalent and is thus used in this study for applying sampling and sensitivity analysis techniques to a more tractable set of input parameters. A domain expert selected what were believed to be the most common variables for a retail building and assigned minimum, maximum, and default values for each of them.

Table 1. Twenty inputs and ranges sampled to generate two sets of 1,024 simulations for sensitivity analysis.

Class	Name	Short Name	Field	Default	Min	Max
Schedule: Compact	CLGSETP_SCH	SC_CL6	Field 6	23	16.1	29.9
	HTGSETP_SCH	SC_HT4	Field 4	16	11.2	20.8
	HTGSETP_SCH	SC_HT6	Field 6	22	15.4	28.6
Lights	Back_Space_Lights	Li_BaSp	Watts per Zone Floor Area	9	6.3	11.7
	Core_Retail_Lights	Li_CoRt	Watts per Zone Floor Area	18.5	12.95	24.05
	Front_Entry_Lights	Li_FrEn	Watts per Zone Floor Area	12	8.4	15.6
	Front_Retail_Lights	Li_FrRt	Watts per Zone Floor Area	18.5	12.95	24.05
	Point_Of_Sale_Lights	Li_POS	Watts per Zone Floor Area	18.5	12.95	24.05
Electric Equipment	BackSpace_MiscPlug	Eq_BaSp	Watts per Zone Floor Area	8.2	5.74	10.66
	CoreRetail_MiscPlug	Eq_CoRt	Watts per Zone Floor Area	3.3	2.31	4.29
	FrontRetail_MiscPlug	Eq_FrRt	Watts per Zone Floor Area	3.3	2.31	4.29
	PointOfSale_MiscPlug	Eq_POS	Watts per Zone Floor Area	22	15.4	28.6
ZoneInfil: FlowRate	Back_Space_Infil	ZF_BaSp	Flow per Ext Surface Area	0.00033	0.00023 1	0.000429
	Front_Entry_Infil	ZF_FrEn	Air Changes per Hour	1.1	0.77	1.43
	Front_Entry_Infil	ZF_FrRtA	Constant Term Coefficient	0	0	1
	Front_Retail_Infil	ZF_FrRtC	Flow per Ext Surface Area	0.00033	0.00023 1	0.000429
	Point_Of_Sale_Infil	ZF_POS	Flow per Ext Surface Area	0.00033	0.00023 1	0.000429
DsgnSpec: OutdrAir	SZ DSOA Back_Space	DS_BaSp	Outdr Airflow per Area	0.0008	0.00056	0.00104
	SZ DSOA Core_Retail	DS_CoRt	Outdr Airflow per Area	0.00175	0.00122 5	0.002275
Sizing: Parameters	Sizing:Parameters	Sz_Heat	Heating Sizing Factor	1.25	0.875	1.625

We define a set of 1,024 simulations using a fractional factorials design sampling where each simulation is defined by the min or max value of each input parameter. The parameter ranges shown here were selected to demonstrate an obvious simulation failure condition in cases where heating setpoint is above cooling setpoint, which happens on 384 (37.5%) of the simulations. This causes EnergyPlus to exit with a severe error "DualSetPointWithDeadBand:

Effective heating set-point higher than effective cooling set-point - increase deadband if using unmixed air model.” It should be noted that simulation failures from specific input combinations are often much more difficult to understand and challenging to anticipate. Even when such combinations are apparent, placing additional rules to sample within the safe state space of input combinations can break some of the statistical advantages of popular sampling methods for experimental design. We demonstrate this by forcing overlapping setpoints to their non-overlapping defaults; this defines a 2nd set of 1,024 simulations with no missing simulation outputs. This comes at the expense of changing the multi-dimensional uniformity of fractional factorials sampling. The impact of this modification will be shown in the resulting analysis.

3.3 Supercomputing

The number of simulations resulting from all the different types of sampling strategies presented above is fairly large. Several supercomputing systems were used for running the simulations and some of the subsequent analysis to tractably execute the simulations in a fair amount of time. The systems included Jaguar/Titan (the fastest supercomputer in the world at the time), Nautilus from the Joint Institute for Computational Sciences, the Oak Ridge Institutional Clusters, and several other supercomputers. For the residential building parameters, a portion of Markov Order 3 simulations were run, but the combinatorial complexity made this intractable to run even with high performance computing resources. The described residential simulations, in addition to simulations of the 3 most popular commercial building types (warehouse, retail, medium office), were calculated and used for multiple data-mining experiments. The Autotune project (Garrett, 2013), in an effort to promote open science, has made a portion of the final 267TB (26.9 trillion data points) of ~8 million E+ simulations data publicly available at <http://rsc.ornl.gov/autotune>.

One of the challenges in using large supercomputing resources was that EnergyPlus is an inherently input-output (I/O) bound simulation engine. While this is not relevant for desktop computing environments, running large batteries of simulations on supercomputing resources greatly increases the number of I/O operations and becomes the biggest bottleneck in the computation. The use of memory mapped file systems local to individual nodes in the supercomputing cluster greatly helped in mitigating the performance issues allowing production level scalability to over 130,000 processors (Sanyal, 2014). To further improve the I/O performance, batches of files were compressed to a single archive. This was done on both the input and the output side, which were then extracted to their target locations. Some of the simulated output was inserted to a MySQL database and setup for querying. The overheads of importing into a database are significant when the number of simulations is large. A large fraction of the simulations were organized into a directory hierarchy determined by simulation type, location, and vintage, and saved as simple flat files. One of the benefits of using statistically designed ensembles is a sharp reduction in I/O requirements.

4. ANALYSIS

4.1 Sensitivity Analysis of ZEBRAlliance Simulation Outputs

To observe the relative change of different outputs and be able to contrast between them, all data ranges were normalized between 0 and 1 and plotted together for comparison. Figure 1 illustrates the sensitivities of all 82 outputs for an ensemble of annual simulations with their normalized means and standard deviations. It is easy to perceive the relative change in the outputs and the parameter combination can be easily looked up.

4.2 Input Parameter Sensitivity Analysis for Total Electricity and Gas Use

The full 2nd set of 1,024 simulations, where some set-points were adjusted to sample within the safe state space, is used here with minimal impact on results as discussed in the following section. We fit a multivariate analysis of variance model, treating the months as repeated measures as we discussed earlier. With 20 factors and all 2-way interactions model fit to 1024 observations, we still have 778 degrees of freedom left for the error term. The multivariate Pillai statistic gives strong significance (below 0.0001) to all factors with the exception of Li_FrEn, ZF_FrEn, and ZF_RtA, which are above 0.5. Among the 380 possible 2-way interactions, 58 are significant at the 0.01 level. While significance indicates sufficient data to measure a non-zero effect (that is, it confirms that we can measure it with our experiment), the actual size of the effect is ultimately what is important. In Figure 2 we show the terms with an effect (coefficient) of over 1,000 KWh energy for Electricity in at least one month. Similarly, we show terms with over 1,000KWh effect (coefficient) for Gas in Figure 3. The factors are labeled with the Short Name, followed by its upper setting. For example, SC_CL6_29.9 is the additive effect (in each month) of the SC_CL6 factor at its 29.9 setting on energy use when all other factors are at their reference setting (lower limit).




Figure 1. Summary of normalized means and standard deviations for all Markov Order 1 annual simulation outputs.

As one would expect, a higher thermostat setting on SC_CL6 (cooling) reduces electricity use. We get more reduction in the summer (bluish points) than in the winter (reddish points). Sz_Heat has a positive effect on electricity use but its interaction with SC_CL6 is negative, indicating that higher thermostat settings reduce the cost of a larger Sz_Heat. Figure 3 shows the corresponding result for Gas. In Gas energy use, it is notable that the “(Intercept)” is missing among the larger effects. This is due to the parameterization and the use of lower levels as the reference for higher levels. Lower reference levels for Gas are very low (or zero) as Gas is used for heating only. This is also reflected in the fact that all main effects are positive, meaning that higher thermostat settings lead to higher energy use.




Figure 2: Factors with largest effects on electricity use. Note the different scales on the bottom plot that would place it in a very narrow range on the top plot. Coefficients are electricity use in KWh.

To illustrate how a fractional factorial data set provides estimates for many settings, we include graphs of estimated monthly mean effect for each factor level with uncertainty in Figure 4. Here, the reference level is the average of the response at the two settings for each factor. This provides a visualization of effects across months of the year. As can be seen from the effect sizes in Figures 2 and 3, some factors have an effect that is one or more orders of magnitude larger than the rest. As a result, we only provide plots for factors where there is a visual difference between the two settings. The same data points are shown in each pair of graphs colored by the levels of a factor. The first two pairs are for the factors with adjusted set-points, resulting in three levels. The estimated line for each factor includes 95 percent uncertainty bounds. Note that each graph contains two lines (three lines in the first two pairs), one for each level of the factor. This illustrates how the effect of each factor can be a simple average because all other factors are balanced. We discuss sensitivity to the adjusted set-points in the next section.

4.3 Sensitivity to Missing Data and Adjusted Set-Points

To address the difference between a clean two-level fractional factorial sample and our data that contained a number of adjusted set-points, we repeated our analysis with those middle settings removed (i.e. missing). This reduces the data to 640, down from 1024. We found that the top main effect and interaction effect estimates (illustrated in

Figures 2 and 3) were minimally affected except that some interactions with them were not estimable. There is enough other data to still have good estimates of the decomposed effects.




Figure 3: Factors with largest effects on gas use. Note the different scale on the bottom plot that would place it in a very narrow range on the top plot. Coefficients are Gas use in KWh.

5. CONCLUSIONS AND FUTURE WORK

In conclusion, we have used three samplings of 156 input parameters for a residential building and shown the impact on a subset of the full 96 EnergyPlus outputs. In addition, we used 20 input parameters on the stand-alone retail reference building to show the limitations and impact of standard experimental designs like fractional factorials on sensitivity results derived from standard analysis of variance techniques. These results should not be extrapolated to all residential or medium office buildings and are shown only to illustrate the impact of sampling methods and practical limitations on the resulting analysis. In addition, we expect the sampling and ensemble analysis at various scales to help us gain insight into unique issues of building energy modeling, especially at different scales of simulation. We also expect the analytic approaches employed for understanding the thermal properties of building envelopes to be beneficial for software calibration and building design.

In future work, we plan to implement *in-situ analysis* (where the analysis runs on the same resources at the simulation) for large-scale simulation data on high performance computing to mitigate I/O bottlenecks. A simulation-agnostic framework has already been established for running hundreds of thousands of individual simulations in parallel. Instead of writing simulation results to disk and later analyzing, we plan to use in-situ methods to perform the analysis while the result files are still in RAMDisk (a portion of RAM which functions as a traditional disk/path but operates at 10-20x traditional disk speed). We will use the pbdR (Programming with Big Data in R: <http://r-pbd.org/>) framework to develop tightly coupled in-situ analysis. Scaling in-situ analysis methods with simulation code, balancing simulation and analysis workload and effective use of resource co-located with the simulation will be critical challenges. When the analysis is complete, only the necessary results can be written to

disk, foregoing long disk write times (~45 minutes for 45TB of simulation data generated in a production-quality 23-minute job on Titan). This is easily scalable for analysis such as finding the absolute minimum or maximum of

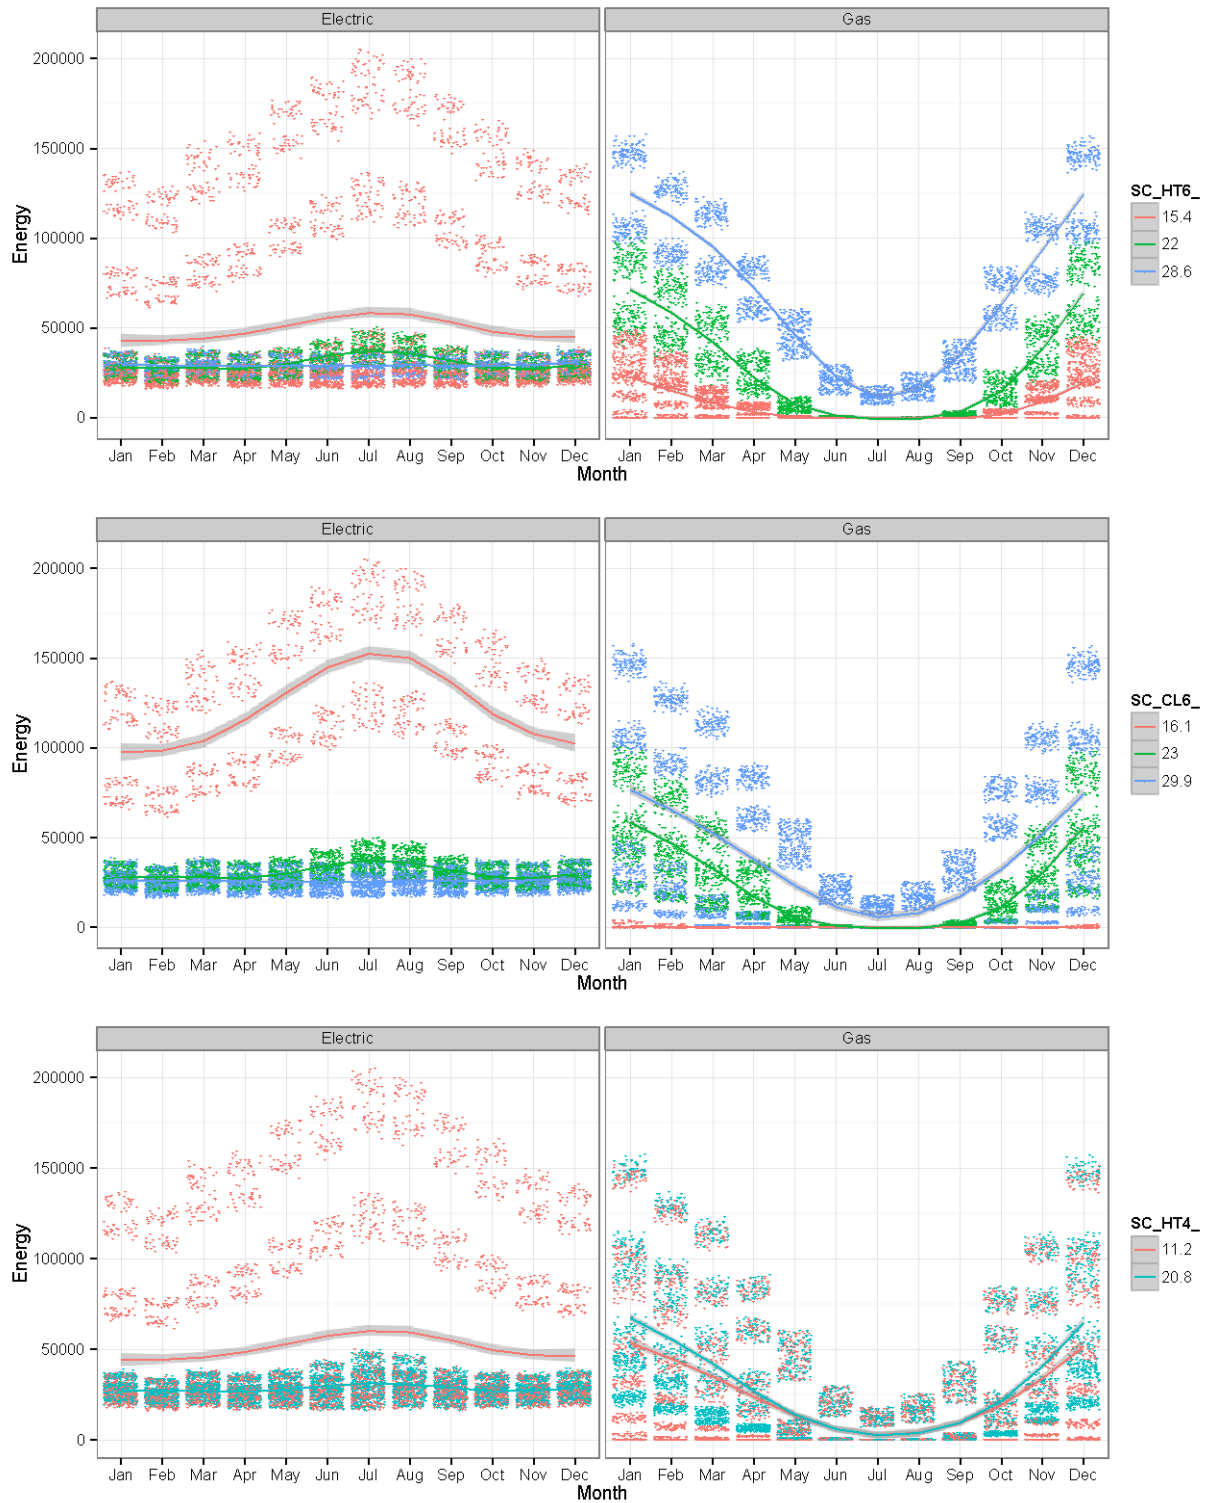


Figure 4: Data plots for monthly Electric and Gas energy use for factors with visually noticeable effect

an output variable, since global extrema (impact to an output variable) can easily be computed from several local lists of extrema. However, a remaining challenge of mining big data is that while ANOVA and similar analyses can trivially be performed on individual blocks of simulations within each node, as demonstrated in this study, merging analysis from multiple sets of simulations across all compute nodes can be challenging for many types of analysis.

REFERENCES

- Biswas, K., Gehl, A., Jackson, R., Boudreaux, P., and Christian, J. 2012. Comparison of Two High-Performance Energy Efficient Homes: Annual Performance Report, December 1, 2010–November 30, 2011. Oak Ridge National Laboratory report ORNL/TM-2011/539. Available: <http://info.ornl.gov/sites/publications/files/Pub34163.pdf>
- Crawley, D.B., Hand, J.W., Kummert, M., and Griffith, B.T., 2008, Contrasting the capabilities of building energy performance simulation programs, *Building and Environment*, vol. 43, p. 661-673.
- Deru, M., Field, K., Studer, D., Benne, K., Griffith, B., Torcellini, P., Liu, B., Halverson, M., Winiarski, D., Rosenberg M., Huang, J., and Crawley, D. 2011. US Department of Energy commercial reference building models of the national building stock. NREL Technical report TP-5500-46861.
- Garrett, A., New, J., and Chandler, T. 2013. Evolutionary Tuning of Building Models to Monthly Electrical Consumption." Technical Paper DE-13-008. In Proceedings of the ASHRAE Annual Conference and ASHRAE Transactions 2013, volume 119, part 2, pp. 89-100, Denver, CO, June 22-26, 2013.
- Hand, D. J. and Taylor, C. C. (1987) *Multivariate Analysis of Variance and Repeated Measures*. Chapman and Hall
- Miller, W.A. et al. 2012. ZEBRAAlliance: An Alliance Maximizing Cost-Effective Energy Efficiency in Buildings. Available: <http://www.zebralliance.org>.
- Mee, R., *A Comprehensive Guide to Factorial Two-Level Experimentation*, a monograph for Springer: New York, July 2009.
- R Development Core Team (2014). R: A language and environment for statistical computing. R Foundation for Statistical Computing, URL <http://www.R-project.org>.
- Reddy, T.A., Maor, I., Jian, S., and Panjapornporn, C. 2006. Procedures for Reconciling Computer-Calculated Results with Measured Energy Data. *ASHRAE Research Project RP-1051*, 2006.
- Sanyal, Jibonananda, New, Joshua R., Edwards, Richard E., and Parker, Lynne E. (2014). "Calibrating Building Energy Models Using Supercomputer Trained Machine Learning Agents." In Proceedings of the Journal on Concurrency and Computation: Practice and Experience, February, 2014
- Ulrike Grömping (2014). R Package FrF2 for Creating and Analyzing Fractional Factorial 2-Level Designs. *Journal of Statistical Software*, 56(1), 1-56. URL <http://www.jstatsoft.org/v56/i01/>

ACKNOWLEDGEMENT

Funding for this project was provided by field work proposal CEBT105 under the Department of Energy Building Technology Activity Number BT0201000. We would like to thank Som Shrestha and Mahabir Bhandari for their selection of residential building parameters, and Buzz Karpay for his selection of the stand-alone retail parameters used in this study. This research used resources of the Oak Ridge Leadership Computing Facility at the Oak Ridge National Laboratory, which is supported by the Office of Science of the U.S. Department of Energy under Contract No. DE-AC05-00OR22725. Our work has been enabled by data analysis and visualization experts from NICS (National Institute for Computational Sciences) through an NSF XSEDE Extended Collaborative Support Service(ECSS) program. We would like to thank XSEDE and the Oak Ridge Leadership Computational Facility for making the systems available.

Oak Ridge National Laboratory is managed by UT-Battelle, LLC, for the U.S. Dept. of Energy under contract DE-AC05-00OR22725. This manuscript has been authored by UT-Battelle, LLC, under Contract Number DE-AC05-00OR22725 with the U.S. Department of Energy. The United States Government retains and the publisher, by accepting the article for publication, acknowledges that the United States Government retains a non-exclusive, paid-up, irrevocable, world-wide license to publish or reproduce the published form of this manuscript, or allow others to do so, for United States Government purposes.