Supercomputer Assisted Generation of Machine Learning Agents for the Calibration of Building Energy Models

Oak Ridge National Laboratory
University of Tennessee in Knoxville

Jibonananda Sanyal, Joshua New, and Richard Edwards

{sanyalj, newjr}@ornl.gov, redwar15@utk.edu

More information and data access at http://autotune.roofcalc.com

Building Energy Modeling and its importance

Building Energy Modeling (BEM) utilizes physics-based models to determine energy consumption for a building. The accurate modeling of buildings can serve as an important tool in determining your energy footprint, designing energy-efficient buildings, or optimizing return-on-investment for existing buildings.

In the year 2010, the ~119 million US buildings consumed 41% of the primary energy (up from 39% in 2006) constituting a larger share of the US energy economy than either transportation (28%) or industry (31%). This is approximately $400 billion (up from $220 billion in 2006) with buildings consuming 74% of all electricity and 34% of all natural gas produced. The Department of Energy's Building Technologies Office has a significant stake in improving the energy footprint and efficiency of the buildings sector for economic, social, and environmental benefits.

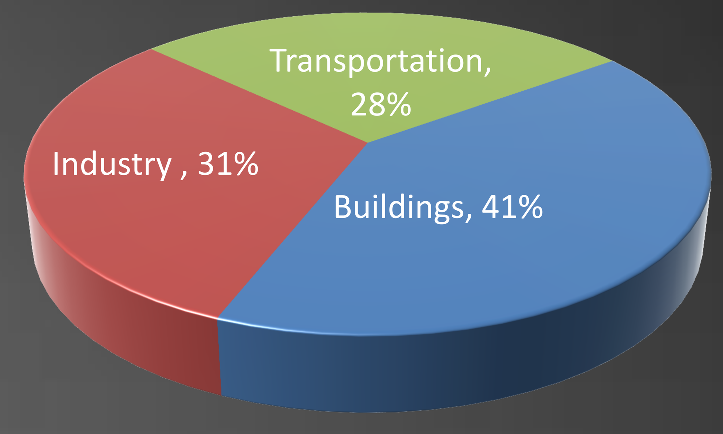

U.S. Primary Energy Consumption in 2010

Business-as-usual method of calibration

In the "business-as-usual" model, experts manually adjust different input parameters until satisfactory results are obtained. The time required for this can sometimes be in the order of a few months! A typical input to the simulation engine can contain over 3,000 input parameters for a regular residential building. Many of these are difficult to obtain without sufficient laboratory test facilities, and some vary after the point of measurement.

Autotune The Autotune methodology

The goal of the "Autotune" project is to save the time and effort spent by energy modelers in adjusting simulation input parameters to match the measured data by providing an 'easy' button. An initial model of the building is provided to Autotune along with a set of measured sensor data. Autotune then spins off the trained machine agents and returns a tuned model of the building allowing energy efficiency analysis at speed and scale.

Traditional model tuning vs. "Autotune"

But, isn't each building unique?

Yes, but buildings must conform to code and DOE has 16 reference buildings that are most representative of the U.S. building stock. These are usually the starting models used energy modelers. Machine learning agents are trained on simulation data from residential, medium office, stand-alone retail, and warehouse buildings, with plans to cover all 16 building types.

The following breakdown is being used:

- Residential, medium office, stand-alone retail, warehouse
- Vintage: Old, Recent, and New constructions
- Across ASHRAE defined 16 climate zones in the U.S.

So, how big is the parameter space?

A typical building could have over 3000 input parameters. Sampling all of this space is neither tractable nor defensible. We've employed the following strategies to trim this space:

- Experts selected 156 of 3000+ input parameter
  - Brute force using 3 levels: 5×10^{52} E+ simulations
  - 3.2×10^{28} lifetimes of the known universe on *Titan*!
- 14 parameter full combinatorial subset of most important ones
- Markov Order 1 and 2 sampling for single- and pair-wise effects
- Latin Hypercube
- Hierarchical groups

These experimental setups over variables including building type, vintage, and climatic zone result in about 8 million simulations:

- Residential: 5 million
- Medium Office: 1 million
- Stand-alone retail: 1 million
- Warehouse: 1 million

For the residential building, we have ground truth data from robotized and heavily instrumented ZEBRAlliance homes

Supercomputing challenge – EnergyPlus is desktop software

We took the 'for-desktop ' software and scaled it to shared-memory *Nautilus* as well as on distributed-memory *Titan*, which is now the world's 2nd fastest supercomputer.

Traditionally, such programs scale poorly. We heavily employ the RAM based file system to mitigate Lustre performance , pack multiple input and output files into tarballs in RAM, carefully balance number of simulations, heap space, and *tmpfs* space on each node, and implement asynchronous mechanisms in an otherwise bulk-synchronous operation.

- We are one of the largest users of XSEDE Nautilus having used over 300,000 SUs on this machine
- For a 131,072 core run on Titan, we generated 45TBs in 68 minutes!
- About 270+TB raw data for the 8 million simulations
- Data transfer: can be cheaper to re-run simulations
- Working with XSEDE /NICS Parallel Big-Data R team for in-stream analysis who scaled R to use 12,000 cores on XSEDE/NICS Kraken
- Working on leveraging XSEDE/SDSC Gordon's Hadoop as a job capability

Nautilus

No of Processors	No of E+ sims	Wall-clock Time (h:mm:ss)	Time/E+ task (mm:ss)
32	32	0:02:08	2:08
64	64	0:03:04	3:04
128	128	0:04:11	4:11
128	1024	0:34:24	4:18
256	2048	1:25:17	10:40
512	1024	0:18:05	9:02

Scalability on Nautilus

Titan

No of Processors	E+ Tasks	Wall-clock Time(mm:ss)
64	256	18:34
128	512	18:22
256	1024	20:30
512	2048	20:43
1024	4096	21:03
2048	8192	21:11
4096	16384	20:00
8192	32768	26:14
16384	65536	26:11
65536	262144	44:52
131072	524288	68:08

Scalability on Titan

Kraken

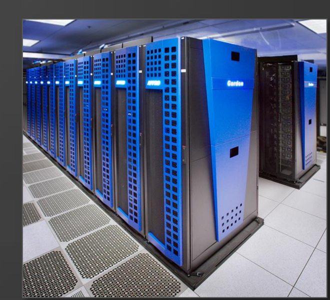

Gordon

Generating Machine Learning Agents

Multiple machine learning algorithms and software implementations were integrated into a single HPC-enabled system known as MLSuite to create "cheat sheets" via big data mining that speed up the Autotune process.

- Support Vector Machines
- Genetic Algorithms
- (Non-)Linear Regression
- Self-Organizing Maps
- C/K-Means
- FF/Recurrent Neural Networks

For commercial buildings		ASHRAE G14 Requires	Autotune Results
Using Monthly utility data	CV(RMSE)	30%	0.318%
	NMBE	10%	0.059%
Using Hourly utility data	CV(RMSE)	15%	0.483%
	NMBE	5%	0.067%

So, how well did we do?

Below **0.5%** error

ASHRAE guidelines require 30% to be legally useful