


International Conference Organised by
IBPSA-Nordic, 13th-14th October 2020,
OsloMet

BuildSIM-Nordic 2020

Selected papers



SINTEF Proceedings

Editors:

Laurent Georges, Matthias Haase, Vojislav Novakovic and Peter G. Schild

BuildSIM-Nordic 2020

Selected papers

International Conference Organised by IBPSA-Nordic,
13th–14th October 2020, OsloMet

SINTEF Academic Press

SINTEF Proceedings no 5

Editors:

Laurent Georges, Matthias Haase, Vojislav Novakovic and Peter G. Schild

BuildSIM-Nordic 2020

Selected papers

International Conference Organised by IBPSA-Nordic,

13th–14th October 2020, OsloMet

Keywords:

Building acoustics, Building Information Modelling (BIM), Building physics, CFD and air flow, Commissioning and control, Daylighting and lighting, Developments in simulation, Education in building performance simulation, Energy storage, Heating, Ventilation and Air Conditioning (HVAC), Human behavior in simulation, Indoor Environmental Quality (IEQ), New software developments, Optimization, Simulation at urban scale, Simulation to support regulations, Simulation vs reality, Solar energy systems, Validation, calibration and uncertainty, Weather data & Climate adaptation, Fenestration (windows & shading), Zero Energy Buildings (ZEB), Emissions and Life Cycle Analysis

Cover illustration: IBPSA-logo

ISSN 2387-4295 (online)

ISBN 978-82-536-1679-7 (pdf)



© The authors

Published by SINTEF Academic Press 2020

This is an open access publication under the CC BY-NC-ND license

(<http://creativecommons.org/licenses/by-nc-nd/4.0/>).

SINTEF Academic Press

Address: Børrestuveien 3

PO Box 124 Blindern

N-0314 OSLO

Tel: +47 40 00 51 00

www.sintef.no/community

www.sintefbok.no

SINTEF Proceedings

SINTEF Proceedings is a serial publication for peer-reviewed conference proceedings on a variety of scientific topics.

The processes of peer-reviewing of papers published in SINTEF Proceedings are administered by the conference organizers and proceedings editors. Detailed procedures will vary according to custom and practice in each scientific community.

Global Marginal Carbon Footprint Evaluation of Internet Services with Building Energy Models

Eric Kumar, Erica Cochran Hameen, Wei Liang
Carnegie Mellon University, Pittsburgh PA
School of Architecture

Abstract

Over the last decade, operational efficiency optimizations in data center (DC) facilities designs have curtailed their absolute power demands. However, their physical footprint is still growing. This growth in capacity may offset the efficiency gains at best or may even increase the power demand in absolute terms. Such offsetting effects make system level life cycle environmental footprint evaluations complex for distributed services operating in these DCs.

To support system level evaluations for services operating in DCs, this research uses simulated network traffic profiles and coincident energy generation source aware building energy models (BEM) to evaluate the marginal carbon footprint of a globally spanning network of data centers. The result of this research is a prototype of a BEM based framework that can be used in environmental decisions to characterize DC life-cycle costs at internet service levels.

Introduction

Globally distributed internet services are ubiquitous today. These services are driving the demand to build mega-watt scale data center facilities (DC) distributed through out the globe. In a 2015 study data centers were forecast to consume as much as 13% of the global energy production by 2030 (Andrae and Edler, 2015). More optimistic models from the US Department of Energy for the US showed a curtailment with up to fifty percent decline in energy compared to the industry's use of 2% of the power produced by the national grid in 2005 as a result of using state of the art efficiency and consolidation practices (Shehabi, 2016). As more and more parts of society are transitioning to data center dependent online paradigms, the absolute demand of data centers is growing. The volume of growth is exemplified by the capital costs being invested across the world in constructing these data centers.

The growth in data center capital construction costs will reach \$89 billion by 2027, a significant increase from the \$45 billion spent in 2018 (Insight-Partners,

2019). Capital costs aren't the only commitments for data center owners however. DC owner's also incur significant operational costs throughout the entire operational lives of their facilities.

Over a 20-year life of a continuously operating data-center facility, the use-phase energy costs can exceed its capital costs while having a much larger ecological footprint. Given the accumulation of costs and impacts over the life of data centers, there is a need for a robust model that couples the information technology (IT) and building systems with their energy supply sources through its useful life. In this article a geographically extensible model that accounts for the workloads, building systems, and power utility grid is presented.

In the rest of this paper a model for coincident energy demand and marginal energy costs (MEC) of a set of data centers is proposed and developed. First, the proposed method uses a hybrid model consisting of EnergyPlus and Python programming modules as developed by the researcher in (Kumar, 2020a). For this article, the researcher's original hybrid model is modified as described in the methodology section to be more indicative of the current operations of cloud data centers. Then second, the resulting time series of operational energy demand profile is used as input to a Python based open source MEC simulation tool that is also extended as described in the methodology section.

The MEC calculations are based on the Dispatch Optimized Systems Cost of Energy (DOSCOE) model developed by Platt (Platt et al., 2017). DOSCOE provides a linear programming platform that computes the monetary costs and several environmental emissions associated with operating a mix of dispatch-able and non-dispatch-able energy sources. In this context, dispatch-able energy sources are those that can be controlled to meet demand. Natural gas power plants are examples of dispatch-able sources of energy, where the plant operators can control the mass flow rates of the combustion gases to curtail generation rates. On the contrary are the non-

dispatchable sources, where the energy generation is dependent on extrinsic factors. An example of non-dispatchable energy source is solar; where cloud cover greatly effects the generation rate and no power can be generated at night.

Modeling the costs for a mixture of dispatchable and non-dispatchable generators is complex as most non-dispatchable sources are only accounted for as opportunistic supply sources when sizing the power generation infrastructure. By design the dispatchable generators are sized to meet the full demand in a worst-case condition when no dispatchable power is available (Platt et al., 2017). However, when non-dispatchable power is opportunistically available, it supplements the grid; allowing dispatchable sources to be turned down to part loads. This saves fuel costs for dispatchable sources (ie natural gas generators), but it leaves their physical infrastructure stranded and underutilized.

This research’s coupling of the DOSCOE model with a BEM quantifies the monetary and ecological costs associated with operating data centers in grids with mixtures of energy sources bound by physical constraints. The developed model is evaluated by matching the marginal costs of energy with data center demands on hourly intervals. Furthermore, the resulting model is inclusive of the stranded costs of the underutilized power infrastructure.

In the next section, Background, informative context is provided to set the proper use-case for this framework. Then in the Similar Works section, past literature which have quantified the ecological life cycle cost of data centers and internet services are summarized. In the Methodology section, the two main modules of the software implementation of this research are presented. Specifically these modules are a modified version of hybrid building energy model and a novel MEC model based on DOSCOE. After the details of the models are presented, the results are discussed in the Results and Discussion section followed by the Conclusions.

Background

Data centers are critical to modern internet experiences for billions of people. They are the key enablers for disseminating information in real time regardless of people’s location. As an example of location agnostic services geographical dispersed infrastructure, Figure 1 illustrates a geographically distributed architecture that enables data centers to provide globally consistent internet experience that have become the status quo over the last decade. Distributed software architectures implemented in the server clusters at the leaf-nodes consist of a collection of autonomous computing resources that appear to its users as a single coherent system (van Steen and Tanenbaum, 2018).




Figure 1: General Topology of data center networks. The 1:1 metropolitan area to building relationships are shown for clarity only. Pink: indicates the first layer - ISP-Metropolitan links (physical). Magenta: indicates the second layer Metropolitan-DC Building with a full mesh connection (physical). Gray: represents the cross-connection between Data Centers (logical). Green/Red: Cluster to Cluster links (logical).

In Figure 1, a data center network stack with three hierarchical WAN layers are shown. The first hierarchical layer, the global level, has a wide area network (WAN) connected to two internet service providers (ISP) as its root node. For the purposes of this work, the WAN is an abstraction of a network that connects a set of data center facilities with each other. The second layer of the global level are the metropolitan regions. In this layer, the ISP links are shown as coming in from top and the outgoing links from the bottom serve local distribution networks to buildings within the metropolitan areas. The final layer of the global level are the data center buildings, where the global network links connect to clusters of servers. Inside each data center there may be another independent hierarchy as shown by the server clusters. The physical network links shown in Figure 1 connect the buildings to each other through the WAN. However, the physical links (pink, magenta) is transparent for the clusters, which directly communicate with each other through logical links (gray, green, red).

Large cloud data center operators have championed WAN systems for global up-time (service availability) and user experience optimizations (Sushant et al., 2013) (Hong and Kandula, 2013). Cloud data centers are a specific class of data centers. They serve business functions that require custom IT and building systems tuned for optimality in total costs of ownership models. These facilities house numerous services that can be controlled by network load balancers, with each data center limited by its physical infrastructure’s capacity. As described by Sushant and Hong, the WAN networks can shift loads between data center on command. The efficacy of a strategy for shifting computational loads depends on the capacity headroom of the physical resources at the data center where the loads are shifted to.




Figure 2: Activity distribution of a sample of 2 clusters, each containing over 20,000 servers, over a period of 3 months. The x-axis indicates the server utilization rates (0 to 1) and the y-axis indicates the fraction of time during the 3 months (Barroso et al., 2018). The figure on the right (b) shows that peak utilization exceeds 75% with little variances over the 3 months, whereas the right figure (a) shows that utilization is below 60% at all times.

The location fungibility enabled by WANs is desirable from a service application performance point of view, as latency can be significantly reduced for specific markets and applications by minimizing the round trip communication times with consumers. Furthermore for cloud DCs, researchers have demonstrated cross awareness of BEMs and network models to optimize DC sustainability (Tripathi et al., 2017; Kiani and Ansari, 2017). Beyond the sustainability gains, there is a potential for meaningful financial incentives by balancing workloads across DC, thus reducing the peak demand on any specific DC and lowering the required physical capacity that must be built.

In the context of BEMs, there is a trade-off with the non-deterministic nature of network load balancing described above. Implementations of load balancing strategies using WAN's make the IT workloads temporally (and geographically) unstable, rendering it elusive for building modelers to reason about. As an example of temporal and geographical instability of data centers, Figure 2 (Barroso et al., 2018) shows the difference in utilization rates for two server clusters from Barroso's operational experience. This behavior of IT loads is something that WAN aware BEMs can help characterize and exploit. One possible means of exploitation is to flatten the peak of each cluster by routing the excessive workloads to another similarly provisioned cluster with lower coincident demands and sufficient building systems capacity headroom. However, building energy simulation need to be aware of more than just energy to make the right environmental decisions. To effectively optimize environmental objectives, BEMs coupled with MEC models can quantify the energy's global warming potential in terms of carbon footprint to shift loads from one building to another.

Based on the literature reviews and references cited, there are no publicly available simulation frameworks that couple the dynamic data center workloads and

the coincident carbon footprint associated with powering their load. The marginal cost of energy coupled building energy model described in this research is the first publicly available tool to allow owners, designers, and researchers to quantify the carbon footprint of data center operations accounting for granular supply and demand matching of power. Next, a literature review of past works concerning the carbon footprint of data-centers and digital services is presented in the Similar Works section.

Similar Works

There are two notable past works that look at the energy footprint for distributed sets of data centers. First, Tripathi considers hardware capital costs alongside with energy acquisition costs to quantify the total costs of ownership in (Tripathi et al., 2017). Tripathi's framework is dynamic in terms of workloads but it is not aware of the building energy dynamics. In the second work, Kiani and Ansari describe a geographical load balancing strategy that exploits green energy mix in the utility grid (Kiani and Ansari, 2017). However, their load balancing scheme doesn't provide insights into how the load balancer evaluates the building energy demands or how the greenness of the energy supply is obtained.

Using a life cycle assessment framework, Whitehead demonstrates a comprehensive data center site level life cycle costs analysis in (Whitehead, 2015). All energy use in Whitehead's models were deduced from annualized PUE values, precluding coincident energy source evaluations with their framework. Similarly the Green Grid's data center life cycle assessment guideline is limited to PUE as their suggested basis for the operational energy proxy (The Green Grid, 2012). The PUE metric is shown in Equation 1, it has proven to be the most popular method for data center operational efficiencies since 2006 (Wikipedia.com, 2020).

$$PUE = \frac{E_{total}}{E_{IT}} \quad (1)$$

E_{total} = Total Power Used at Facility

E_{IT} = Power Consumed by IT Equipment

Other's have evaluated the costs of internet services (Taylor and Koomey, 2008; Shehabi et al., 2014). Taylor and Koomey quantify the energy and greenhouse gas implications of online Advertising circa 2008. While Shehabi evaluates the energy and greenhouse gas implications of video stream circa 2014 (Shehabi et al., 2014). Together these works provide a taxonomy that can be followed to assess internet service costs in terms of energy use.



Figure 3: Data center locations

Methodology

In this section the MEC model's coupling with the BEM is described. The resulting model maintains a strict partition between the two technical domains. The first part of the model simulates the hourly energy demands of a set of five data centers in EnergyPlus (EP) using the method demonstrated in (Kumar, 2020a). These data centers are distributed across the globe as shown in Figure 3. Then in the second part, the data center demands are matched with the respective region's utility power generation sources for the coincident hour to assess the MEC during that hour. These parts are described in the following subsections.

A fundamental component of the demonstrated implementation of this framework is the network traffic simulator from (Kumar, 2020a). The simulation produces a time-series profile of the network traffic that a hypothetical service will get at a particular data center site. The hypothetical services are based on Wikipedia and segregated according to the natural language of the Wikipedia pages. By segregating services by natural languages allows the treatment of each language as a 'service platform'. With these network simulations, one or more languages can be supported from a single data center with the constraint that the sum of the IT workloads does not exceed the capacity of the data center facility.

Building Energy Model

For the first part, a sufficient building energy demand profile is simulated as in (Kumar, 2020a). The IT workload profile characterizing the power and cooling load is obtained by simulating the 50th quantile of daily network traffic to each data center using Kumar's method from (Kumar, 2020b). There are three notable differences between Kumar's original BEM and the model presented here. The specific changes are:

1. Setting $2kW/m^2$ as the IT equipment load density to represent cloud DCs with hot-aisle/cold aisle containment resulted in run-time errors. This persistently led to thermal runaway condi-

tions for the Singapore and San Francisco sites in the new model. In order to keep the building envelope form-factor and the construction materials the same, this simulation's IT power load density is changed to $1.0kW/m^2$.

2. In these simulations, the data center model has air distribution flow control with approach temperatures specified. Flow control with approach temperature method calculates the temperature differences between the IT hardware boundaries and the air handling equipment. This method is more representative of modern data-center operations and allows modeling ASHRAE 90.4 Standard's requirement of hot-aisle / cold-aisle compartmentalization. The alternate method in EnergyPlus considers the entire data center room as a well-mixed environment, consistent with the modeling from (Kumar, 2020a). While using the approach temperature method, the cooling set-point is changed to 27 degrees Celsius from 25 degrees used in (Kumar, 2020a). This 2-degree adjustment corresponds with the approach temperature between the air handling equipment discharge and the inlets of the ITE; as there is no mixing of the supply air before it enters the ITE.
3. As the third and final change, the load distribution values in the EnergyPlus input file (IDF) were revised from the SequentialLoad setting to UniformPLR. In the former, equipment is activated in the order listed in the IDF. With this specification each piece of equipment ramps sequentially from its idle state to full capacity, before subsequent equipment is enabled. In the latter load distribution specification, all equipment are loaded in parallel to each other. Another available setting for load distribution, the Optimal specification, was also tried for this field, but it crashed the simulations.

To validate that the proposed building energy model configurations with the above changes produce reasonable results, each data-center is simulated with two models. In the first model, the economizer for the direct evaporative cooler (DEC) limits are increased while in the second model the default settings are maintained. The summary of the changes are indicated in Table 1, while the comparison of total site energy between the two models for each set of the simulations is illustrated in Figure 4.

In Figure 4, the pairs of simulations for a DC site are presented. In the figure, the red bar indicates the total site energy of the model with economizer limits increased and the IT loads reset by the network coefficient (RL). The blue bar indicates the default economizer values (NoRL) from Kumar and with statically configured IT load profiles (ie cyclic day/time load). From the figure it can be observed that the increased



Figure 4: Total Site annual energy for statically (NoRL) configured IT load and dynamically reset (RL) IT loads.

economizer settings lead to more energy use than the default values for three of the five DCs. This increase in energy demand is attributed to the load resetting, where the network aware model increased the IT load for more hours of the year.

Table 1: Economizer settings for the two models

IDF Object	Variable	Default	Increased
DC-OA	Econ. Max db-C	23	28
DC DEC	Evap. Max db-C	20	28

The next subsection introduces the marginal costs of energy model, which will take the results from the BEM model described above.

Regional Marginal Costs of Energy

The second module requires a time-series profile of the electrical grid’s power source attributions, where the intervals of the power generation values match the BEM energy profile. The corresponding site and regional grid model pairs are indicated in Table 2. For this model the energy generation regional grid profiles are obtained from Platt, who provides profiles for 13 US regions in DOSCOE [5]. Singapore and Amsterdam data-centers grid profiles are constructed as described below.

Table 2: Data Center Site and Power Grid

DC ID	DC Location	Regional Grid
SFO	San Francisco, CA	California
CAR	Carrollton, Texas	ERCOT
ASH	Ashburn, VA	Midatlantic
AMS	Amsterdam	Netherlands
SIN	Singapore	Singapore

Each U.S region’s grid profile is composed of hourly demands and corresponding production capacity of several power generation technologies. Hourly values for each region’s demand, solar, wind, coal, coal with

cryogenic capture, coal amine gas scrubbing, and nuclear are defined. These grid regions are representative of three out of five of the data-center locations being simulated.

For Singapore, the International Energy Agency (IEA) provides a top down view of the annual energy production from various technologies. The IEA data shows that the renewable penetration in the energy supply for Singapore accounted for only 1.6% of the total energy demand in 2016 (IEA, 2017). Due to this negligible contribution from renewable sources and lack of hourly generation data, it is assumed that all demand is met by natural gas power generators, consistent with other sources (USEIA, 2016b). Figure 5 indicates the energy production capacity and the associated carbon emissions for each of the utility grid regions. The demand profile for Singapore in 2016 is obtained from (Singapore-Government, 2016).

For the Amsterdam data center, the energy profile of Netherlands is used. The IEA indicates that in 2018, 12% of Netherlands’ energy demands was met by renewable sources. This is a meaningful contribution from renewable, therefore a more granular generation profile is prudent as opposed to Singapore which did not require evaluations of the energy mix at any given time. To formulate a more granular resolution of the Netherlands’ generation, the IEA data is supplemented with data from the Open-Power-System-Data to characterize the time series profile of renewable energy (Open-Power-System-Data, 2019). OPSD provides hourly data for the renewable sources only. The balance of the energy demands in the Netherlands is met by fossil fuel-based generators, namely 35% by coal and 42% by natural gas (USEIA, 2016a). The DOSCOE grid profiles don’t have any corresponding field for bio-mass, so the bio-mass generation indicated in OPSD is lumped in with the fossil fuel generators. In the discussion section, some validation for this approach is presented. The Netherlands also produces nuclear energy, but only the annual production rates were obtained (USEIA, 2016a). In this works implementation of the DOSCOE, the annual nuclear production for Netherlands is distributed equally over the year and modeled as a constant (non-dispatchable) base load throughout the year.

The MEC coupled BEM algorithm is given below. In the algorithm two inputs are required. The first is DOSCOE[region], it is a two-dimensional array formatted as described in (Platt et al., 2017). It indicates hourly grid profiles of the power demand and power capacity of the available energy sources at the corresponding hour. The second input, traffic.language, is a one-dimensional vector indicating the network traffic to the respective site from Kumar (2020b). In the algorithm, for each language, it’s traffic to the respective data center is checked. If there are traces of a language routed to a site, the algo-

rithm performs the BEM simulation by invoking EnergyPlus. This resulting demands from EnergyPlus are then added to the region’s grid demand profile. If a language does not have any traffic to a particular site than, the data center site does not do any work and the BEM simulation is bypassed.

Algorithm 1 MEC coupled BEM algorithm

Require: DOSCOE[region] & traffic.language[site]
for site and region in DC.site and DC.region **do**
 if traffic.language[site].all != 0: **then**
 $DEMAND_{DC} \leftarrow BEM(site, traffic.language)$
 $DOSCOE[region].demand += demand[site]$
 end if
end for
 $CO_2 \text{ footprint} = GridSim(region, rps)$

Where $DEMAND_{DC}$ is the marginal demand the data center puts on the power grid,

The second step of algorithm quantifies the marginal carbon footprint of the grid with the added data center loads by running DOSCOE’s Grid Simulator. In this step, the renewable portfolio standard (rps) argument specifies the percentage of renewable energy mix for the region.

In the next section the results from the methodology are discussed.

Results & Discussions

The resulting values of the carbon footprint for each data center is summarized in Table 3. The energy model for the 491 kW data center has been scaled by 1000 to represent the metro scale of data centers impact of the regional grid. The specific data center model represents only a small fraction of hyper-scale data centers campuses that reach 500-Mega-Watts. The table further indicates the carbon footprint of each of the languages being served from the data centers. The values are indicative of the marginal CO_2 emitted by adding the data center demand to the respective grid.

Table 3 indicates the total marginal carbon footprint for each language in the last row by summing each of the language columns. While the total marginal carbon footprint of each data center is obtained by summing the rows as shown in the last column. The English pages have the most traffic globally and as expected English has the largest marginal carbon footprint due to its data-center operations. From the physical data center perspective, the model shows that the Ashburn data center has the highest carbon footprint. This can be attributed in part to the energy source’s carbon footprint relative to other locations as shown in Figure 5.

Figure 6 is a histogram plot of the PUEs for all the data centers in the evaluated network. Each color in the figure represents a specific data center. The

PUE is indicative of the utilities system efficiencies; namely the cooling and the power distribution. The x-axis in Figure 6 indicates the PUE value and the y-axis indicates the number of hours the data centers have at specific values. As a concrete example, the bars representing Netherlands site is labeled as amsr_l.PUE and has most hours at or below a PUE of 1.1. Other data centers also have a PUE of 1.1. Stacking the curves from the different data centers indicates the global network’s time spent at a specific PUE. A major influencing factor for the PUE is the local weather conditions.

An understanding of the cooling design weather conditions may help reason about the differences in PUE. Lower PUEs generally coincide with cooler weather, which leads to more favorable thermo-mechanical equipment operating points. In practical terms, cooler temperatures enable facilities to be cooled with economizers for more hours and they also lower the equipment’s energy demand when economizers are not fully operational. For the set of the data center analyzed in this work, the ASHRAE cooling design conditions are shown in Table 4 (from ASHRAE. (2009)).

As an implementation detail, the DOSCOE model proved to be quite sensitive to the data structure of the load profiles. For example any null value in the profile resulted in breaking the linear solver. Also, solar and wind energy are required to be non-zero for at least one hour of the year. This requirement is a practical constraint, but in the profile developed for this work the values are arbitrarily set to a low value across all hours. Specifically, the defaults setting used in the work is 0.1 MW for each hour.

Another implementation measure required the discounting of facility’s base loads. Without this discount, each model accounted not only the incremen-

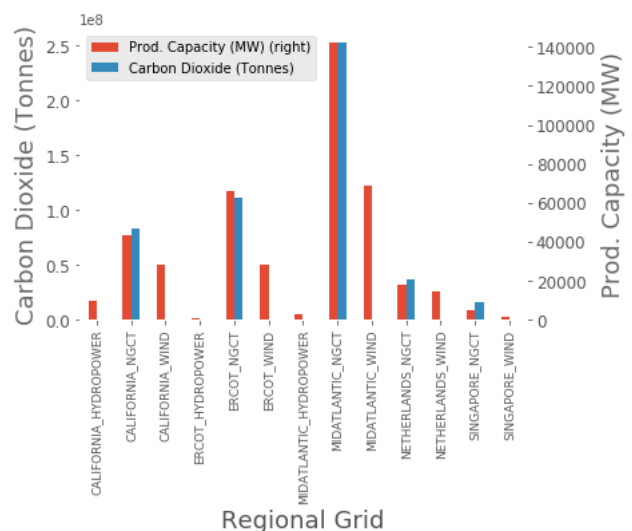


Figure 5: Energy Product Capacity and Carbon Emissions.

Site	German	English	Spanish	French	Japanese	Russian	Chinese	Total
California	0.0	1.24x10 ⁶	0.0	0.0	0.0	0.0	0.0	1.24x10 ⁶
Carrollton	0.0	6.59x10 ⁴	1.20x10 ⁶	0.0	0.0	0.0	0.0	1.26x10 ⁶
Ashburn	0.0	4.38x10 ⁶	8.0x10 ⁵	1.17x10 ⁵	0.0	0.0	0.0	1.36x10 ⁶
Amsterdam	2.81x10 ⁵	2.25x10 ⁵	9.85x10 ⁴	2.67x10 ⁵	0.0	3.15x10 ⁵	4.67x10 ⁴	1.23x10 ⁶
Singapore	0.0	3.74x10 ⁵	0.0	3.71x10 ⁴	3.35x10 ⁵	3.91x10 ⁵	7.87x10 ⁴	1.22x10 ⁶
Total	2.81x10 ⁵	2.35x10 ⁶	2.09x10 ⁶	4.21x10 ⁵	3.34x10 ⁵	7.06x10 ⁵	1.25x5	6.31x10 ⁶

Table 3: Data Center Operation Carbon Footprint (Tonnes of CO₂ Emitted over the Year)




Figure 6: Stacked histogram of PUE from all Data Centers.

DC	0.4%DB	0.4%WB	1%DB	1%WB	2%DB	2%WB
Amsterdam	82	67.5	78.1	66.2	74.6	64.4
Ashburn	93.5	75.1	90.8	74.3	88.2	73
Singapore	91.7	79.4	91.1	79.4	89.9	87.2
Carrollton	100.4	74.5	98.4	74.6	96.2	74.8
San Francisco	83	63	78.3	62.1	74.4	61.2

Table 4: ASHRAE Annual Cooling and Dehumidification Design Conditions in °F ASHRAE. (2009)

tal load placed on the DC from a service’s incoming traffic, but also it would account for the entire baseline loads in each EnergyPlus run. For sites with multiple services, this results in significantly over-accounting of the base loads, as it would be included for all the services in the DC. To address this issue, a post processing step apportioned the baseline load to each service/language by using the traffic coefficients. For the attributions, an initial condition for the traffic coefficient is set to 0. With this initial condition, the BEM provides the baseline energy required by the data center (i.e. all equipment operating at their lowest turn-down rates). The traffic coefficient scaled this value, bound by the constraint that the sum of the coefficients is strictly equal to 1 for all time-steps of the simulation.

Conclusion

The network dependencies between physically dispersed data center resources make system level building design decisions a challenge to reason about. This research has presented a quantitative method to couple the network dependencies of data centers with

their building energy and grid level marginal costs of energy. The research extended the PUE metric with utility grids to indicate carbon footprint of data centers. Furthermore, as shown in Table 4 the presented method allows analysis of data center systems at global service levels.

As pointed out in the background section, there is a lack of bottoms up building design and carbon footprint models. The modular methodology of this research is a novel means of coupling abstracted service level metrics (WAN traffic in this case) with physics based building energy simulation (EnergyPlus in this case). This integrated tool set can be used by data center system modelers to optimize their deployment across geographical bounds. Future work should consider controlling the network traffic based on building energy simulations to achieve global optimality in real time operational environments.

References

- Andrae, A. and T. Edler (2015). On global electricity usage of communication technology: Trends to 2030. In *Challenges*, Volume 6, pp. 117–157.
- ASHRAE. (2009). *ASHRAE Handbook: Fundamentals*. Atlanta, GA: ASHRAE.
- Barroso, L. A., U. Hozle, and P. Ranganathan (2018). *The Datacenter as a Computer: Designing Warehouse-Scale Machines 3rd ed.* Morgan and Claypool.
- Hong, C.-Y. and S. Kandula (2013). Achieving high utilization with software-driven wan. In *SIGCOMM’13 ACM, Hong Kong*.
- IEA (2017). International energy data and statistics. Retrieved: March 03, 2020.
- Insight-Partners (2019). Data center construction market. Retrieved: May 12, 2020. <https://www.theinsightpartners.com/reports/data-center-construction-market>.
- Kiani, A. and N. Ansari (2017). On the fundamental energy trade-offs of geographical load balancing. Volume 0163-6804/17, pp. 170–175.
- Kumar, E. (2020a). Towards energy simulations for proportionally designed and controlled data centers. In *Accepted for Building Simulation and Opti-*

- mization 2020. Loughborough (UK), 21-23 September 2020.
- Kumar, E. (2020b). Wide area network based data center energy simulations for internet services. In *Accepted for 11th International Conference on Improving Energy Efficiency in Commercial Buildings and Smart Communities*. Frankfurt (Germany), October, 2020.
- Open-Power-System-Data (2019). Time series open-power-system-data. Retrieved: April 02, 2020. https://data.openpowersystem-data.org/time_series/20190605.
- Platt, J., O. Pritchard, and D. Bryant (2017). Data center construction market 2019-2024 — global industry overview by size, share, future demand, latest research by competitors, segmentation and regional forecast. *SSRN 08*.
- Shehabi, A. (2016). United States Data Center Energy Usage, United States Department of Energy.
- Shehabi, A., B. Walker, and E. Masanet (2014, may). The energy and greenhouse-gas implications of internet video streaming in the united states. *Environmental Research Letters 9*(5).
- (2016). *Half Hourly System Demand*. Retrieved: January 02, 2020. <https://data.gov.sg/dataset>.
- Sushant, J., S. Mandal, J. Ong, A. Singh, U. Holze, and A. Vahdat (2013). B4: Experience with a globally-deployed software defined WAN. In *SIGCOMM'13 ACM, Hong Kong*.
- Taylor, C. and J. Koomey (2008). Estimating the use and greenhouse emissions of internet advertising, IMC squared.
- The Green Grid, T. (2012). Data center life cycle assessment guidelines. Retrieved: May 13, 2020. <https://www.thegreengrid.org/en/resources/library-and-tools/236-Data-Center-Life-Cycle-Assessment-Guidelines>.
- Tripathi, R., S. Vignesh, V. Tamarapalli, and D. Medhi (2017). Cost efficient design of fault tolerant geo distributed datacenters. In *IEEE Transactions on Network and Service Management*.
- USEIA (2016a). U.S. Energy Information Administration. Retrieved: May 02, 2020. www.eia.gov/international/analysis/country/NLD.
- USEIA (2016b). U.S. Energy Information Administration. Retrieved: May 05, 2020. www.eia.gov/international/analysis/country/SGP.
- van Steen, M. and A. S. Tanenbaum (2018). *Distributed Systems Third Edition*. Switzerland: Maarten van Steen.
- Whitehead, B. (2015). The life cycle assessment of a UK data centre. In *International Journal of Life Cycle Assessments*, Volume 20, pp. 332–349.
- Wikipedia.com (2020). Power usage effectiveness. Retrieved: August 01, 2020.