



This is an electronic reprint of the original article. This reprint may differ from the original in pagination and typographic detail.

Siltala, Mikko; Brannvall, Rickard; Gustafsson, Jonas; Zhou, Quan Physical and Data-Driven Models for Edge Data Center Cooling System

Published in: 2020 Swedish Workshop on Data Science, SweDS 2020

DOI: 10.1109/SweDS51247.2020.9275588

Published: 29/10/2020

Document Version Peer-reviewed accepted author manuscript, also known as Final accepted manuscript or Post-print

Please cite the original version:

Siltala, M., Brannvall, R., Gustafsson, J., & Zhou, Q. (2020). Physical and Data-Driven Models for Edge Data Center Cooling System. In *2020 Swedish Workshop on Data Science, SweDS 2020* Article 9275588 IEEE. https://doi.org/10.1109/SweDS51247.2020.9275588

This material is protected by copyright and other intellectual property rights, and duplication or sale of all or part of any of the repository collections is not permitted, except that material may be duplicated by you for your research use or educational purposes in electronic or print form. You must obtain permission for any other use. Electronic or print copies may not be offered, whether for sale or otherwise to anyone who is not an authorised user.

© 2020 IEEE. This is the author's version of an article that has been published by IEEE. Personal use of this material is permitted. Permission from IEEE must be obtained for all other uses, in any current or future media, including reprinting/republishing this material for advertising or promotional purposes, creating new collective works, for resale or redistribution to servers or lists, or reuse of any copyrighted component of this work in other works.

## PHYSICAL AND DATA-DRIVEN MODELS FOR EDGE DATA CENTER COOLING SYSTEM

Mikko Siltala<sup>\*†</sup>, Rickard Brännvall<sup>\*‡</sup>, Jonas Gustafsson<sup>\*</sup>

RISE Research Institutes of Sweden, Department of Computer Science Luleå, Sweden

## Quan Zhou

# Aalto University, Department of Electrical Engineering and Automation, Espoo, Finland

### ABSTRACT

Edge data centers are expected to become prevalent providing low latency computing power for 5G mobile and IoT applications. This article develops two models for the complete cooling system of an edge data center: one model based on the laws of thermodynamics and one data-driven model based on LSTM neural networks. The models are validated against an actual edge data center experimental set-up showing root mean squared errors (RMSE) for most individual components below 1 °C over a simulation period of approximately 10 hours; which compares favourably to state-of-the-art models.

*Index Terms*— Edge, Data center, LSTM, Cooling System, Thermal Energy Storage

#### **1. INTRODUCTION**

In recent years, data center (DC) energy efficiency has received increasing attention due to the rapid increase in the number of DCs and their energy usage worldwide [13]. In a data center, it is common for approximately a third of the energy usage to be spent on cooling [2], which is essentially wasted energy. Cooling systems are often only static or reactive, attempting to hold a constant air temperature in the DC. Further optimization could be achieved using proactive control algorithms, such as the model predictive control (MPC) method. These methods use mathematical models to predict the response of the system, but it is challenging to create accurate models of the thermodynamics of a DC cooling system.

The state-of-the-art methods in mathematical modeling of data center cooling system are physical modeling, Computational Fluid Dynamics (CFD) modeling [11, 20], and recently also data-driven modeling including neural networks [1, 6, 8]. CFD models target the air flow and heat transfer in data centers, but usually only for the air interactions in the server room. Physical models make broad simplifications in describing the dynamics of the cooling system, such as with the chiller and dry cooler, in for example [3, 21], or with the chilled-water heat exchanger [18], compact models that additionally include the thermal mass for room, plenum, walls, floor, ceiling and a water storage tank [9], or use networks of thermal nodes to model the internal dynamics of servers [14]. Moreover, neural networks have been used to directly predict the effectiveness of DC cooling systems [6, 16], and more rarely the system thermal responses [1].

An edge data center provides cloud computing and storage services at the edge of a network, for example at the crossing point between a mobile network and the wired internet. It is expected that edge data centers will be widely distributed and located close to the users to reduce the latency for mobile/internet-of-things (IoT) applications that rely on computation/storage outside of the devices. For easier integration in urban environments and with power distribution systems they may be equipped with batteries, renewable energy sources and thermal energy storage [4].

A motivation for this work is to investigate how to construct accurate models for small edge data centers with thermal energy storage using the existing modeling approaches - where such accurate models can give insight to the internal states of the entire cooling system and contribute to the design of more efficient facilities and control systems.

*Our contribution* is a direct comparison of a physical model and a data-driven model for the cooling systems of a small data center equipped with thermal energy storage.

## 2. EXPERIMENTAL SET-UP

The EDGE data center is a small data center laboratory built at the RISE ICE data center research facilities at Luleå, Sweden, which functions as a test bed for the edge data center

<sup>\*</sup>Work funded by Vinnova through ITEA 3 project no. 17002, AutoDC.

<sup>&</sup>lt;sup>†</sup>Corresponding author: mikko.siltala@ri.se, carried out this study at RISE for his Master's at school of Electrical Engineering, Aalto University

<sup>&</sup>lt;sup>‡</sup>Corresponding author: rickard.brannvall@ri.se, at RISE, is also pursuing PhD studies at EISLAB, Luleå University of Technology, Sweden



**Fig. 1**. EDGE lab: (1) The data center module, a container for the data center. (2) Two server racks. (3) CRAH (mounted above the racks). (4) The coolant water storage tank. (5) Chiller. (6) Coolant pipes going to the roof. (7) Cooling tower (on the roof). (8) Batteries. (9) The measurement and control system. (10) Picture of the solar panels (on the roof). (11) Microgrid inverter.

concept. In addition to 10 kW of IT, the RISE EDGE lab is equipped with a thermal energy storage (TES) tank and a micro-grid with photo-voltaic cells and batteries allowing experimentation on using alternative energy sources, reliability and partial self-sufficiency for micro-grid connected data centers. We use an experimental set-up reported earlier [4], which is briefly summarized below. Figure 1 shows the experimental set-up where the main components are labeled in the caption. The cooling system consists of components 1 to 7.

- 1. *Data center module*. A sheet metal container, which is separated into a cold and hot aisle (full containment).
- 2. *Two server racks* are positioned between the aisles with 38 Dell PowerEdge R430 servers each. The servers have a total idle load of about 9 kW, and a maximum load of roughly 11.5 kW.
- 3. An air-to-water heat exchanger sits above the racks providing cooling air in the closed loop between the hot and the cold aisle.
- 4. *The storage tank* has a capacity of 2000 liters and is filled with water that is circulated through the module heat exchanger and chiller. Flow diffusers and nozzles (red cones in picture) help reduce internal currents in the tank.
- 5. *The chiller unit* has three modes of operation: the chiller mode, the free cooling mode, and a partial free cooling mode. The first two modes cool the incoming water by vapor-compression technology or liquid-to-liquid heat-exchange technology, respectively, and the third mode uses both methods. The free cooling mode is used when the outside temperature is low and vapor-compressor is not required to increase the temperature difference between the water and glycol.
- 6. *Cooling tower pipes* carry the glycol-mixture from the chiller to the drycooler and back again.

- 7. *The cooling tower* is an air-liquid dry-cooler, where outside air is forced through the heat exchanger, removing some of the heat in the liquid received from the chiller.
- 8. *Lead-acid batteries* with capacity to hold 30 kWh of electrical energy.
- 9. A measurement and control system collects sensor data and sends it downstream for storage, visualisation, analysis and input to control algorithms.
- 10. *The solar panels* on the roof of the facility have a maximum capacity of 10 kW.
- 11. *The microgrid inverter* connects the DC system with the external power grid, solar panels and batteries.

The inputs used to control the system are:

- *IT load setpoint* determines the synthetic load placed on servers by the stress-ng software.
- *Chiller output setpoint* is the temperature of the water the chiller outputs to the storage tank.
- *Pump setpoint* controls the pump between the storage tank and module heat exchanger.
- Fan setpoint for the module heat exchanger fan.

Data from the sensors in the hardware equipment as well as data hosted natively by the chiller and microgrid controller are sampled approximately every 30 seconds by a common data collection system (built on the Zabbix software) before being passed into a time series database (KairosDB) for longterm storage that allows access by the visualization (Grafana API) and analysis (Jupyter/Python or Matlab) tools. The data collection solution used in the lab is designed for scalability and is described in previous work [7].

Experiments were conducted to create the training and validation data for the models by running predetermined sequences that alter set-points over time. The first four experiments changed only one set-point in isolation while the other inputs were kept constant. Three more complex experiments



Fig. 2. The EDGE cooling system.

allowed all of the inputs to change throughout a day: from these one was used as validation data, while the other two together with the first four were used for model training. Details for the seven experiment sequences can be found in [17].

#### 2.1. Physical Model

The physical model represents each of the system components as its own model, which are then joined as is done in the block diagram of Figure 3 to create the complete system model. This method of dividing the model to sections is widely used in state-of-the-art solutions [3, 5, 19, 21], and often implemented in Matlab Simulink for simulations and analysis [3, 21]. The method assumes ideal heat transfer, and uses natural constants, nameplate values and qualified estimates based on system dimensions for model parameters such as surface areas, volumes, and equipment weights.

The complete description of the model does not fit in this article (with only the equations running into several pages). Therefore we only state the equations for the data center hot aisle dynamics as an example. For the other components in the heat rejection system we only present a brief summary of the modeling and assumptions - for detail we refer to the original thesis work by Siltala [17].

Data center module. The energy balance equations for the hot aisle and cold aisle are based on the simplified steady-flow equation, the convection heat transfer rate equation and the internal thermal energy change rate equation. The net thermal energy outflow rate is equal to the convection losses through the walls to the ambient air, and the change in the internal thermal energy. Additionally, the IT equipment is assumed to input thermal energy to the hot aisle by the first law of thermodynamics. We then have for the hot aisle

$$\rho_{air}V_{hot}c_p^{air}\frac{dT_{hot}}{dt} = q_{IT} + \dot{m}_{module}^{air}c_p^{air}(T_{cold} - T_{hot}) - h_{hot}A_{hot}(T_{hot} - T_{ambient}) - h_{IT}A_{IT}(T_{hot} - T_{IT})$$

where the first term on the right hand side describes the IT heat injection, followed by the thermal energy flow, and convective and conductive heat transfer. Model parameters like surface area,  $A_{hot}$ , and volume,  $V_{hot}$ , of the hot aisle are estimated, while the specific heat capacity,  $c_p^{air}$ , air density,  $\rho_{air}$ , and heat transfer coefficient over a thin wall,  $h_{hot}$ , are taken from physical handbook tables.

There is a corresponding differential equation for the cold aisle temperature,  $T_{cold}$ , but without the thermal interactions with the IT equipment. Forward Euler discretization is used to approximate the differential equations for numerical simulation. For the heat transfer between the equipment and the hot aisle air we use the lumped heat capacity equation

$$m_{IT}c_p^{IT}\frac{dT_{IT}}{dt} = h_{IT}A_{IT}(T_{hot} - T_{IT})$$

where  $A_{IT}$  and  $m_{IT}$  are the estimated effective surface area and mass of the IT equipment,  $c_p^{IT}$  the specific heat capacity of the servers (assuming steel), and  $h_{IT}$  the heat transfer coefficient into air.

*Module heat exchanger.* We assume that there is negligible heat transfer between the heat exchanger and its surroundings, negligible potential and kinetic energy changes, there is no phase change and the specific heats are constant [10], and use the steady-flow thermal equation for both of the two media, where the net rate of thermal energy removed from the air equals to the net rate of thermal energy output to the water. The Logarithmic Mean Temperature Difference (LMTD) equations complete the description of the counter-flow heat exchanger.

The value of the heat exchanger thermal conductance is unknown and nonlinearly affected by the air and water flow rates and temperatures. Therefore a polynomial of degree 2 in the pump set-point and fan set-point are fitted to experimental data, later to be used as input to the simulations. The dependence on changing temperature is assumed to be weak and neglected in the regression.

Thermal energy storage. Since the storage can have a



Fig. 3. The Edge data center cooling system block diagram.

temperature difference between the top and bottom, the tank is horizontally divided into two imaginary sections along the middle, and the heat flows are calculated for both. The energy balance equations then assume that the temperature is uniform in these two sections.

*Chiller.* Two sets of equations are required to represent the dynamics of the chiller, since it has two operating modes: chiller mode and free cooling mode (see also greyed area in Figure 2). In free cooling mode the exchange of heat happens at a water-to-glycol heat exchanger, described by the LMTD equations similarly to the module heat exchanger. In chiller mode the refrigeration loop is defined by the coefficient of performance (COP), the thermal energy flow and the lumped heat capacity.

*Cooling tower.* The last component of the heat rejection system is the dry cooler on the roof, where the heat is expelled from the system. This component is again a heat exchanger, and the dynamics are solved again using the LMTD equations assuming a counterflow liquid-to-air heat-exchanger, where the heat radiated to the outside air equal to the heat removed from the glycol.

*Mass flow rates.* There are five mass flow rates that are required to be solved for the system as visualised in Figure 3. Two of these were assumed constant based on observations, and three assumed variable: 1) data center module air-flow rate, which depends on the fan set-point, 2) storage to heat exchange water-flow, which depends on the pump set-point, and 3) chiller to tower glycol-flow in free cooling mode, depending on the chiller control signal regulating a 3-way valve. For these, the functional relations were determined by polynomial regression (details in [17]).

*Transport delays.* Estimating the transport delays in the system can be calculated for the sections where fluid is pumped through pipes by comparing the volume of the pipe and the flow rate.

#### 2.2. Data-driven model

The Long Short-Term Memory (LSTM) neural network architecture which can model time series with correlation over long time spans was chosen for the data-driven model [15], because such dependencies and time delays were observed in the system.

*Modular architecture.* The data-driven models should attempt to model the individual components of the cooling system, as well as the entire system. The system is divided into five sections as is illustrated in Figure 3, where the IT module is considered jointly as one component. Mass flows are not modeled explicitly. The outputs of one model are therefore used as inputs for other models, creating a dependency between the models. The model components have two types of inputs: inputs which are not dependent on the other models, such as the system set-points, and inputs which are the output of some other model components. The training of the models is therefore possible separately.

*Initialization.* The internal state of the model can not be directly inferred from observation data. Therefore the initialization is provided by running the simulation with an initialization sequence before predictions can be made.

*Training data.* The training data was chosen so that it contains all but one of the experiments that were conducted with the system. The remaining experiment was used as the validation data.

*Neural network configuration.* The first layer of the network is the input layer, followed by a number of hidden LSTM cells, connected to a dense-layer that predicts the output. The tested networks had 1, 2 or 3 LSTM layers, with number of neurons being powers of two between 2 and 1024.

*Training procedure.* The temperature prediction error was used in the Mean Squared Error (MSE) loss function. The ADAM optimiser by Kingma and Ba [12] was used for loss minimization to find the model parameters. Models were trained for a hundred epochs to determine the best network configuration; the 2-layered models had the best results, and a layer size of 32 provided a high speed per epoch joined with good performance. Further models were trained with 2 layers of 32 neurons each (with about 13000 trainable parameters).

The final models were trained for up to a thousand epochs (roughly 15 hours) to find further performance increases. Early stopping was practised, by which the epoch with the best validation MSE were saved to be analyzed further.

#### 3. RESULTS

Simulation results were obtained in four sets: both for the physical model and the data-driven model, and for each conducting the simulation both jointly and for each component in isolation. Naturally all these can not fit into the article format; instead only a summary of the simulation results are presented for the entire system. For visual validation, we also include example plots of simulated time series for the hot aisle temperatures.

RMSE (°C)	Physical		Data-driven	
	Comp.	Join.	Comp.	Join.
Hot aisle	0.43	1.68	1.72	2.30
Cold aisle	0.19	1.70	0.48	1.72
HE to storage	0.14	1.76	0.39	1.82
Storage to HE	0.66	1.70	0.63	1.55
Storage to chiller	0.39	1.27	0.35	1.09
Chiller to storage	1.93	2.64	1.20	2.55
Chiller to tower	2.13	3.25	0.97	4.31
Tower to chiller	3.55	3.89	1.46	4.08
Mean	1.18	2.24	0.90	2.43

Table 1. Validation RMSE from simulations.

*Physical model validation.* As can be seen from Table 1 (leftmost column), the RMSE of the hot aisle temperature is quite low for the individual component of the physical model, which is also reflected in the first panel in Figure 4 that tells that this model simulation agrees quite well with observations. The largest discrepancy occurs towards the end of the time series, when the fan set-point is changed, and the temperature change is almost instant without thermal lag. A possible cause is that the heat transfer coefficients are assumed static, but in reality they are dependent on the flow rate of air and the temperature difference.

For the joint simulation, illustrated in the second panel of Figure 4, only the initial temperatures are provided to the physical model that is now tasked to simulate all the components together based on the recorded sequence of set-point changes, ambient and outside temperatures. Poorer performance is expected as errors of each component get propagated back to the system, and indeed Table 1 shows that the RMSE is larger for the joint simulation than for its components.

*Data-driven model.* The third panel of Figure 4 with hot aisle temperatures shows that the module component takes roughly 1.5 hours to fully converge, after which the estimate is fairly accurate for the first half of the simulation, and then deteriorates towards the long prediction horizon.

In the joint data-driven model (fourth panel on Figure 5) the first hour of data is used as the initialization period, and during the rest of the simulation the estimates are fed back to the models. The errors for hot aisle temperatures are somewhat larger for the same reasons as with the physical model.

#### 4. ANALYSIS

Judging by the performance metrics in Table 1, overall the physical models are more often the more accurate model, with lower RMSE than those of the data-driven model.

The whole row for tower-to-chiller temperatures in Table 1 stands out by showing the largest inaccuracies overall, e.g. with a RMSE of 3.55 for the individual physical model component. After the experiments it was discovered that the temperature sensor often receives too high temperature measurements during the morning, when the sun shines on the side of the cooling tower where the sensor is located. For comparison, when a temperature time series recorded by the Swedish Meteorological and Hydrological Institute (SMHI) was used instead, the RMSE drops to 0.85 for the same model.

The physical models outperform their data-driven counterparts for the server and heat exchanger components. For the storage tank on the other hand, accuracies are very close for the two models, with a slight edge to the data-driven approach according to the RMSE in Table 1. The chiller and cooling tower however were more accurately modeled by the data-driven approach. These observations indicate that components which can be represented using thermodynamic equations can perform well, while components with more complex dynamics that are difficult to capture in simple equations can benefit from a data-driven approach.

A closer look at the simulation time series for the remaining components<sup>1</sup> show that the data-driven model fails to capture the mode-switching behaviour of the chiller unit. These errors, as well as the cooling tower measurement error with the outside air temperature, appear to propagate to the other components and slowly increase the errors for the entire simulation.

The mean average errors (MAE) reported for state-of-theart physical models is 1.1 °C [21] or 1.7 - 0.5 °C [3] for the cold aisle temperature, and 1.15 °C [21] for the hot aisle temperature. The present work gives MAE of 1.45 °C for the cold aisle, and 1.35 °C for the hot aisle with the joint physical model, and with the joint data-driven model 1.44 °C and 1.84 °C, respectively.

One should note that although the simulation horizons are comparable, the modeled set-ups are not identical, and the present experimental set-up cycles a number of different setpoints, while only IT load is changed in the previous work used for comparison [3, 21]. Additionally, the state-of-the-art models do not simulate the entire cooling system or assume some ideal behavior for components further from the servers. A direct comparison of performance scores is therefore not possible, but ,in comparison with the state-of-the-art models, the present models can give insight to the internal states of the entire cooling system.

A state-of-the-art artificial neural network (ANN) model that estimates the cold aisle temperature for individual discon-

<sup>&</sup>lt;sup>1</sup>not included for space constraints; ref to Fig 29 and 30 in [17]



**Fig. 4**. Hot aisle temperature as estimated individually by the physical model (panel 1), by the joined physical model (panel 2), individually by the server component of the data-driven model (panel 3), and by the joined data-driven model (panel 4).

nected points in time reported a MAE of  $0.6 \,^{\circ}$ C [1]. However, when used to predict longer time series, the performance deteriorated severely, and the estimation error grew to over 5  $^{\circ}$ C in under 300 seconds. The LSTM based ANN solution presented here has better performance over prediction horizons of multiple hours (MAE 1.44  $^{\circ}$ C after 9.75 hours), and therefore outperforms in comparison.

The quality of the data-driven model estimate decreases faster than the physical model estimate (Figure 4). Therefore the physical model can be safely used to simulate the real system for longer time horizons without risking overheating of the real system. However, based on these results both models could be safely used as a plant model in a control system as the "accurate" estimation horizon consists of the estimate settling to several accurate steady states following several setpoint changes, and the simulation would be restarted with new initial data to reset the error.

## 5. CONCLUSION

Two models for an edge data center cooling system were presented: a physical model implemented in Matlab/Simulink, and a data-driven model based on LSTM neural networks implemented in TensorFlow/Keras. The simulations were validated against an actual edge DC experimental set-up showing root mean squared errors (RMSE) for most individual components below 1 °C over a simulation period of 9.75 hours. There were however challenges in correctly simulating the multiple operating modes of the chiller component using both methods. In addition, it was noticed that an incorrectly installed temperature sensor decreased the accuracy of the cooling tower models.

When simulating the entire system using these components, an error in any of their outputs will be propagated to the other components and lower the accuracy of all outputs, which was especially troubling for the errors caused by the chiller and tower components. Overall the physical model achieves a higher accuracy than the data-driven model.

Beyond fixing the outside temperature sensor, future work could involve mixing both physical and data-driven components in a joint grey-box model, to investigate if such an approach can keep simulation within long term physical bounds while allowing fine tuning to observed transient behaviour.

### 6. ACKNOWLEDGEMENTS

This project was funded by Vinnova through ITEA 3 project no. 17002, AutoDC. Thanks to Jeffrey Sarkinen, Mattias Vesterlund and Erik Lundmark at RISE ICE Datacenter for their help in making this work possible.

# References

- J. Athavale, Y. Joshi, and M. Yoda. Artificial neural network based prediction of temperature and flow profile in data centers. In 2018 17th IEEE Intersociety conf on Thermal and Thermomechanical Phenomena in Electronic Systems (ITherm), pages 871–880. IEEE, 2018.
- [2] Maria Avgerinou, Paolo Bertoldi, and Luca Castellazzi. Trends in data centre energy consumption under the European code of conduct for data centre energy efficiency. *Energies*, 10(10), 2017.
- [3] Y. Berezovskaya, A. Mousavi, V. Vyatkin, and X. Zhang. Smart distribution of IT load in energy efficient data centers with focus on cooling systems. In *IECON 2018 - 44th Annual Conference of the IEEE Industrial Electronics Society*, pages 4907–4912. IEEE, 2018.
- [4] Rickard Brännvall, Mikko Siltala, Jonas Gustafsson, Jeffrey Sarkinen, Mattias Vesterlund, and Jon Summers. Edge: Microgrid data center with mixed energy storage. In Proceedings of the Eleventh ACM International Conference on Future Energy Systems, e-Energy '20, page 466–473, New York, NY, USA, 2020. Association for Computing Machinery. ISBN 9781450380096.
- [5] T. J. Breen, E. J. Walsh, J. Punch, A. J. Shah, and C. E. Bash. From chip to cooling tower data center modeling: Part I Influence of server inlet temperature and temperature rise across cabinet. In 2010 12th IEEE Intersociety Conference on Thermal and Thermomechanical Phenomena in Electronic Systems. IEEE, 2010.
- [6] Jim Gao. Machine learning applications for data center optimization. Google AI Research Publication, 2014. [Cited 28.11.2019]. Available at: https://ai. google/research/pubs/pub42542.
- [7] Jonas Gustafsson, Sebastian Fredriksson, Magnus Nilsson-Mäki, Daniel Olsson, Jeffrey Sarkinen, Henrik Niska, Nicolas Seyvet, Tor Minde, and Jonathan Summers. A demonstration of monitoring and measuring data centers for energy efficiency using opensource tools. pages 506–512, 06 2018.
- [8] B. Hadid, S. Lecoeuche, D. Gille, and C. Labarre. Energy efficiency of data centers: A data-driven modelbased approach. In 2016 IEEE International Energy Conference (ENERGYCON). IEEE, 2016.
- [9] C.M. Healey, J.W. VanGilder, M. Condor, and W. Tian. Transient data center temperatures after a primary power outage. pages 865–870, 2018.
- [10] Frank P. Incropera, David P. DeWitt, Theodore L. Bergman, and Adrienne S. Lavine. *Principles of Heat*

and Mass Transfer; International Student Version. Wiley, 7th edition, 2013.

- [11] M. Iyengar, H. Hamann, R. R. Schmidt, and J. Vangilder. Comparison between numerical and experimental temperature distributions in a small data center test cell. In 2007 Proceedings of the ASME InterPack Conference, IPACK 2007, volume 1, pages 819–826, 2007.
- [12] Diederik P. Kingma and Jimmy Ba. Adam: A method for stochastic optimization. In *ICLR*, 2015. URL http://arxiv.org/abs/1412.6980.
- [13] Jonathan G. Koomey. Growth in data center electricity use 2005 to 2010. Technical report, Analytics Press, 2011. [Cited 6.12.2019]. Available at: https: //www.koomey.com/research.html.
- [14] R. Lucchese and A. Johansson. On energy efficient flow provisioning in air-cooled data servers. *Control Engineering Practice*, 89:103–112, 2019.
- [15] J. Schmidhuber. Deep learning in neural networks: An overview. *Neural Networks*, 61:85–117, 2015.
- [16] H. Shoukourian, T. Wilde, D. Labrenz, and A. Bode. Using machine learning for data center cooling infrastructure efficiency prediction. In 2017 IEEE International Parallel and Distributed Processing Symposium Workshops (IPDPSW), pages 954–963. IEEE, 2017.
- [17] Mikko Siltala. Simulating data center cooling systems: data-driven and physical modeling methods. Master's thesis, Aalto University, 2020.
- [18] James W. VanGilder, Christopher M. Healey, Michael Condor, Wei Tian, and Quentin Menusier. A compact cooling-system model for transient data center simulations. In 2018 17th IEEE Intersociety Conference on Thermal and Thermomechanical Phenomena in Electronic Systems (ITherm). IEEE, may 2018.
- [19] E. J. Walsh, T. J. Breen, J. Punch, A. J. Shah, and C. E. Bash. From chip to cooling tower data center modeling: Part II Influence of chip temperature control philosophy. In 2010 12th IEEE Intersociety Conference on Thermal and Thermomechanical Phenomena in Electronic Systems. IEEE, June 2010.
- [20] E Wibron, Anna-Lena Ljung, and T Lundström. Computational fluid dynamics modeling and validating experiments of airflow in a data center. *Energies*, 11, 02 2018.
- [21] G. Zhabelova, M. Vesterlund, S. Eschmann, Y. Berezovskaya, V. Vyatkin, and D. Flieller. A comprehensive model of data center: From CPU to cooling tower. *IEEE Access*, 6:61254–61266, 2018.