Energy Proceedings
Vol 46, 2024
ISSN 2004-2965

Optimizing Waste Heat Utilization from Small-scale Data Centers in Multi-Energy Building Systems

Bürgler Florian[1], Humbert Gabriele[1*], Heer Philipp [1], Koirala Binod Prasad [1]

1 Urban Energy Systems Laboratory, Empa - Swiss Federal Laboratories for Materials Science and Technology, 8600 Dübendorf, Switzerland

(*Corresponding Author: gabriele.humbert@empa.ch)

ABSTRACT

Power consumption from data centers is envisioned to drastically increase in future energy systems. A large share of this power consumption will be, inevitably, converted into heat, proving a huge potential for waste heat recovery opportunities. This work investigates the optimal integration of waste heat recovery from small-scale edge data centers into buildings. A case study using the real-data from NEST building at Dübendorf, Switzerland, is conducted. The results indicate that smaller data centers exhibit better cost optimization. Within the boundary conditions of NEST case study, the self-consumed heat fraction decreases with increasing data center size, and the influence of daily flexibility on heat consumption is negligible.

Keywords: waste heat recovery, data center, optimization, cost minimization, emission minimization, buildings.

NONMENCLATURE

Abbreviations	
WHR	Waste Heat Recovery
DC	Data Center
Symbols	
AC	Annualized cost [€]
AEm	Annualized Emissions [kg_{CO2}/y]
SHF	Self-consumed heat fraction [-]
RR	Required Revenues to break even [CHF/kWh]
E	Energy [kWh]
P	Power [kW]
Q	Heat [kW]

1. INTRODUCTION

The electricity consumption from data centers (DC) is foreseen to increase from 286 TWh in 2016 to 321 TWh in 2030 [1]. Around 52% of the electricity consumed by DCs is used in the IT equipment itself, the rest powers lights, cooling equipment and other supporting systems [2]. Since most of the electrical power consumed in the IT equipment is converted into heat, this offers a huge potential for waste heat recovery [2]. Most studies done on this topic focus on data centers of a medium or large scale, that provides heat to a district heating network or a large facility nearby [3].

In contrast, this work focuses on waste heat recovery on a small scale in edge data centers. Small-scale data centers are considered to have a maximum electric capacity below 250 kW, as defined in the framework of the EcoQube project [4]. Instead of being located in a centralized facility, these small-scale DCs are built 'on the edge' of the network, close to the user. The advantages of edge data centers are their proximity to the end user, which generally results in lower latency [5]. If several edge data centers are managed under the same controller, higher reliability and flexibility of the system can be guaranteed with the possibility of load distribution over the data centers [5]. As an additional benefit, the recovered waste heat can be used directly on-site, thus reducing potential transportation costs and associated heat losses.

However, recovering waste heat from data centers has its constraints and limitations, as the temperature at which it can be utilized is limited by the thermal tolerance levels of the IT components and the chosen cooling technology [6]. In the case of air cooling technologies, which are installed in almost 90% of cases [7], a maximum outlet temperature of 35°C can be achieved [8]. Consequently, domestic heating

This is a paper for 15th International Conference on Applied Energy (ICAE2023), Dec. 3-7, 2023, Doha, Qatar.

applications can be targeted and, with the aid of auxiliary heat pumps, domestic hot water can also be produced.

This research endeavours to delve deeper into the integration of small-scale edge data centers into building energy systems, promoting local utilization of recovered heat instead of channelling it to district heating networks. Although the potential for waste heat recovery from data centers is evident, the literature has not thoroughly explored the holistic study of waste heat recovery and data center operation in multi-energy building systems.

To address this gap, an optimization approach was employed to guide the design of a multi-energy building system incorporating a data center. The study investigates the impact of key design parameters on total cost and generated emissions, specifically focusing on data center size and data center operation flexibility. Through this exploration, valuable insights are gained, offering a comprehensive understanding of waste heat recovery from data centers and its seamless integration into building energy systems.

2. METHODOLOGY

The data center and its integration into a building energy system were modelled in the Ehub optimization tool [9]. This is an open-source multi-objective optimization tool relying on mixed integer linear programming (MILP) that can be used during the planning phases of the multi-energy system design process.

2.1 Ehub optimization tool

In the Ehub tool, five different blocks can be used to define a multi-energy system; namely inputs, outputs, conversion technologies, storage and network technologies. Techno-economic parameters are then provided to each block type as constant values or time series. The design variables considered in the optimization problem were operational and sizing variables for each of the considered conversion and storage technologies. Two different objective functions were formulated:

- *Cost optimization*: *min (C_{tot})*
  where C_{tot} refers to the total cost, calculated as the sum of the annualized investment and O&M costs for all conversion and storage technologies under the imposed input and output constraints;

- *Emission optimization*: *min (Em_{tot})*
  where Em_{tot} refers to the total emissions, calculated as the sum of the embedded and operational emissions for all conversion and storage technologies under the imposed input and output constraints;

Additional information regarding the Ehub tool can be found at [9], while the details for the data center modelling in the Ehub tool are provided in the next section. Specific boundary conditions were imposed deriving from the selected case study of NEST [10], as discussed in section 3.

2.2 Data center modelling

The data center was modelled using various conversion and storage blocks to investigate the DC workload flexibility. Two workload types were defined: *must-run*, representing immediate processing, and *plannable* workload, which can be queued for later. An example of plannable workloads is numerical simulations, which can be scheduled during the day. Both workload types were modelled as demands in the Ehub-tool, as depicted in Fig. 1. To account for operational flexibility, a virtual energy carrier was used as input and demand and workload as output. Plannable tasks were virtually stored, simulating their flexibility, with a custom rule resetting the storage state every 24 hours.

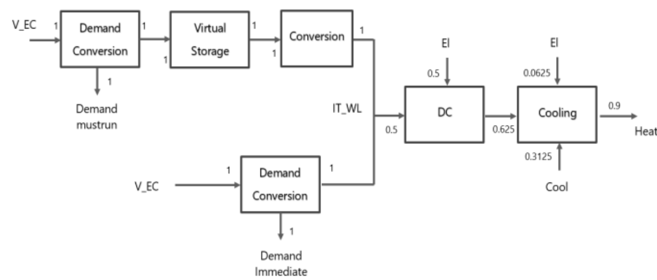

Fig. 1 Flowchart of the data center modelling in Ehub tool

Both workloads feed into the data center conversion block, generating waste heat. This waste heat was subsequently converted into a heat stream in the cooling block. Real measurements from a data center in Sweden were used for workload profiles [11], and different ratios of plannable workload and DC sizes were investigated.

2.3 Key performance indicators

While the optimization problems were solved in agreement with the two objective functions presented in section 2.1, key performance indicators (KPIs) were quantified to assess the benefits of waste heat recovery from data centers. The KPIs are summarized in Table 1. Here, the subscript *ref* indicates a reference case with components' capacities set as the current scenario for the selected case study of NEST [10]. The self-consumed

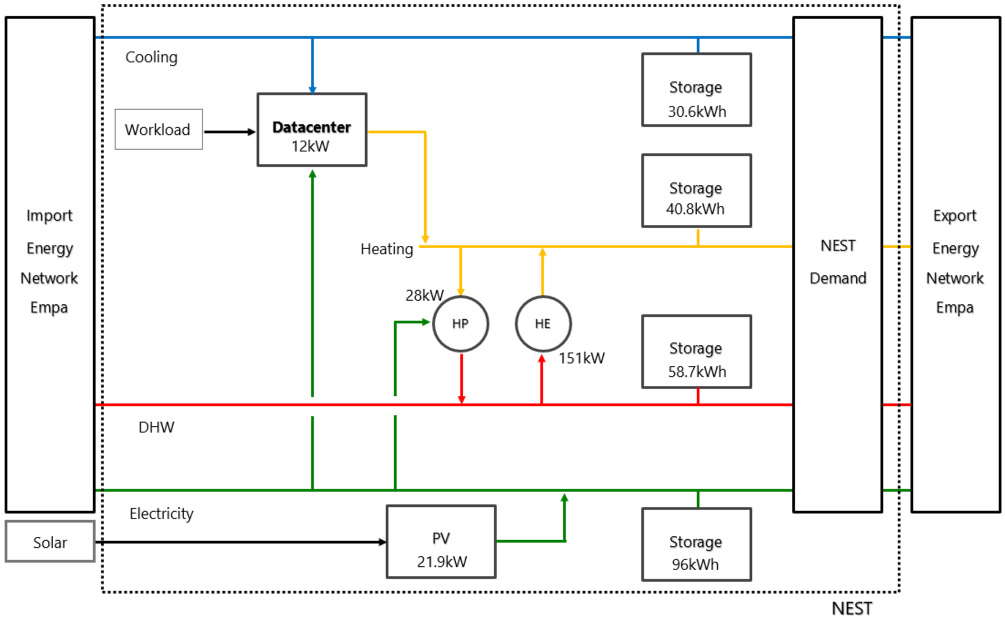

Fig. 2. Schematic of the energy system at NEST

heat fraction, SHF, was defined as the ratio between the heat produced by the DC consumed on-site over the total heat produced by the DC. Finally, the required revenue to break even, RR, was defined as the ratio between the annualized investment and O&M costs attributed to the DC over the energy consumed by the DC.

Table 1 List of selected KPIs.

KPI	unit	Calculation
Annualized Cost: AC	CHF/y,	$C_{tot} - C_{tot,ref}$
Annualized Emissions: AEm	kg_{CO2}/y	$Em_{tot} - Em_{tot,ref}$
Self-consumed Heat Fraction: SHF	-	$\frac{Q_{DC,cons}}{Q_{DC}}$
Req. Revenue to break even: RR	CHF/kWh	$\frac{\mathrm{C_{inv,a,DC}} + \mathrm{C_{O\&M,a,DC}}}{\mathrm{E_{,DC}}}$

A sensitivity analysis for the following key design parameters was conducted:

(i) *DC size*: the data center size was varied in the range from 0 kW, i.e. no data center installed, to 250 kW;

(ii) *IT workload flexibility:* the IT workload flexibility was varied by considering an increasing share of plannable ITworkload from 0% to 100% with a step of 25%;

3. CASE STUDY

Real demand profiles from the NEST building, in Dübendorf, Switzerland, were adopted [10]. NEST is a research and innovation building designed to test new technologies and materials under realistic conditions. The building comprises a backbone that supplies electricity, heating, cooling, and wastewater connections to different modules. Each module is equipped with sensors to dynamically measure demand and production.

NEST's energy system consists of four carriers: electricity, domestic hot water, water for space heating, and cooling. The building connects to the Empa-wide district energy network, allowing heat conversion from domestic hot water to space heating using a heat exchanger with a capacity of 151 kW. The schematic for the NEST building is shown in Fig. 2. A heat pump is also employed to transfer energy from the medium-temperature heating networks (35°C) to the domestic hot water network (65°C) if surplus heat is generated. Each carrier has an integrated storage unit.

Fig. 3. Picture of the small-scale (12 kW) edge data center installed in NEST within the framework of Eco-Qube project [4]

The NEST building consists of various rooftop and façade photovoltaic units for power generation. To simplify the model, a fixed PV profile aggregated all generations from the installed units. NEST has space for additional PV modules on its roof, with the Ehub-tool calculating optimal deployment based on available space. A 12 kW data center with air-cooling unit is currently installed in the NEST as part of the ECO-Qube project [4], as depicted in Fig. 3.

4. RESULTS

First, a base case is discussed that replicates the characteristics of the NEST building. The data center size is set to 12 kW, with assumed 25% plannable workload flexibility. The dynamic profiles for the thermal power are presented in Fig. 4. Here, representative weeks for summer (11th to 17th of July 2022) and winter (10th to 16th of January 2022) seasons are reported. The heat production is represented on the positive y-axis, whereas the negative axis displays the heat consumption or outflow of the system. Notably, there is a significant contrast in heat demand between summer and winter.
During summer, the demand is minimal, typically lasting only a few hours each day, while in winter, heat is required throughout the entire day. Consequently, in winter, all the recovered heat from the data center is utilized, while in summer, the data center can cover all the heat demand, leading to excess heat being exported to the district heating network of Empa. The presence of heat storage allows for the smoothing out of peak loads, although this effect becomes more apparent when the demand roughly aligns with the data center's production. The KPIs values for the base case are reported in Table 2.

Table 2 KPIs values for the base case

	AC [CHF/y]	AE [kg_{CO2}/y]	SHF [-]	RR [CHF/kWh]
Cost opt.	20518	4549	0.632	0.443
Emission opt.	36692	-2880	0.743	0.791

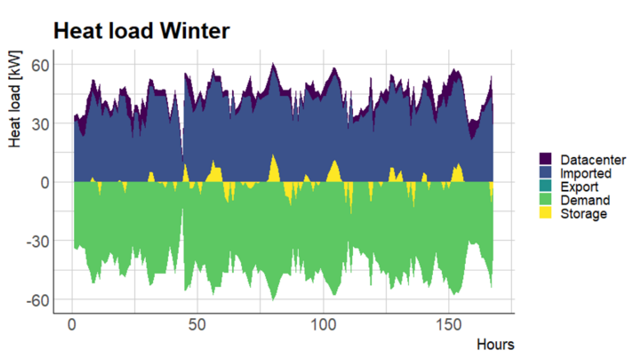

(a)

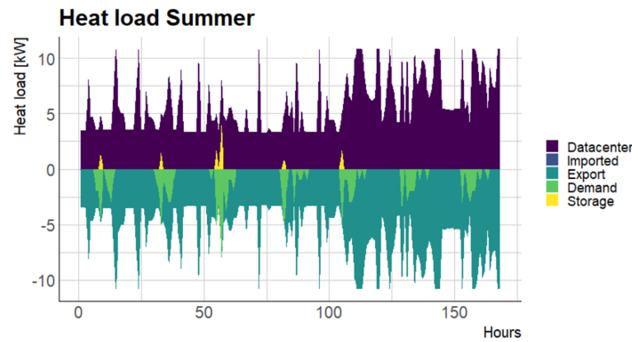

(b)

Fig. 4 Heat consumption and supply profiles of NEST building for representative weeks in (a) winter and (b) summer

4.1 Influence of DC size

The self-consumed heat fraction versus the data center size is shown in Fig. 5 (a). The impact of building demand saturation is evident. With smaller data center sizes, a significant portion of waste heat is utilized within the building. However, as the data center size increases, this fraction diminishes, dropping to less than 15% at a data center size of 250 kW, indicating that a majority of the waste heat is exported rather than being used locally.

The required revenue (RR) per electricity spent versus the data center size is illustrated in Fig. 5 (b). In the case of cost optimization, smaller data centers result in lower RR values, while larger data centers lead to increased costs. This can be attributed to the electricity savings from additional PV installations and the cost savings derived from the exported recovered heat. However, the impact of cost savings from heat recovery diminishes for larger data center sizes as the building's demand is satisfied.

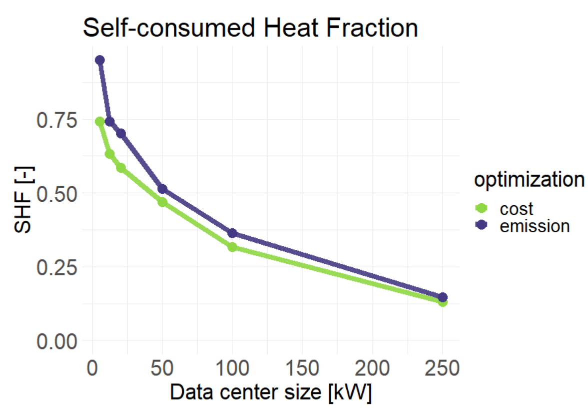

(a)

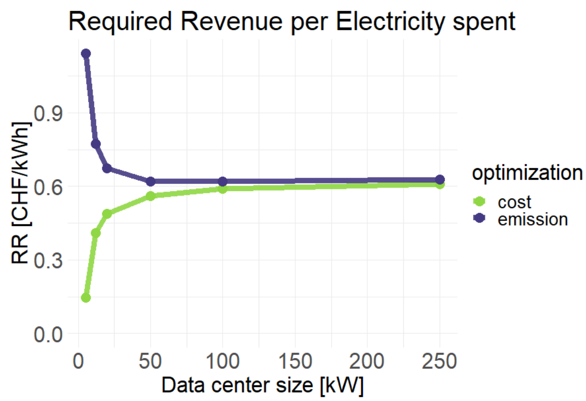

(b)

Fig. 5 KPIs values versus data center size for (a) self-consumed heat fraction and (b) required revenues per electricity spent

Conversely, for emission optimization, the RR values are higher for smaller data center sizes and decrease as the data center capacity grows. This trend approaches the cost optimization results at a data center size of 250 kW. The emission optimization strategy limits PV deployment to ensure all electricity is consumed within the building, thereby underutilizing the potential of additional PV at low data center sizes.

4.2 Influence of IT workload flexibility

The self-consumed heat fraction versus the share of flexible workload for a data center size of 12 kW is shown in Fig. 6 (a). Overall, negligible changes are predicted. Operating the data center with flexibility throughout the day does not result in substantial increases in on-site heat consumption. This outcome is attributed to the seasonal nature of heat demand. In essence, regardless of the daily operational profile, all generated heat is fully utilized during winter, while its consumption is only marginal during summer, mainly due to the low or negligible demand for domestic heating.

Fig. 6 (b) reports about the required revenue per electricity spent versus the flexible workload. Similarly to the previous case, the data center size was fixed at 12 kW. For cost optimization, as the share of the plannable workload increases, the required revenue shows a slight decrease. This can be attributed to the fact that higher flexibility allows the data center to adjust its operations to lower-cost periods, reducing overall expenses. However, the decrease in required revenue is not significant and mainly observed in the first 25% flexibility, indicating that the cost benefits diminish with further increases in flexibility.

On the other hand, for emission optimization, the required revenue trend behaves in the opposite manner. As the plannable workload increases, the required revenue shows a slight increase. This is because higher flexibility allows the data center to optimize its operations to reduce emissions, which may involve higher costs, such as using cleaner energy sources or more energy-efficient technologies.

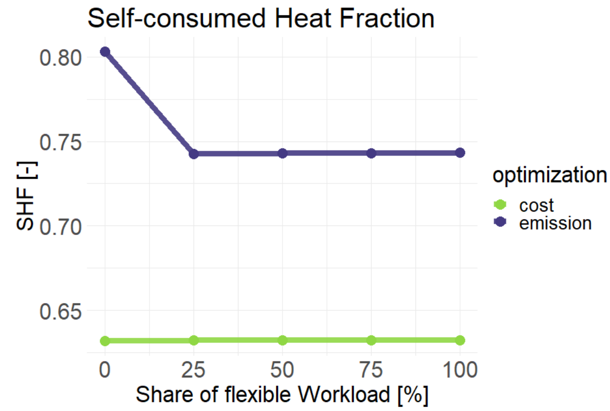

(a)

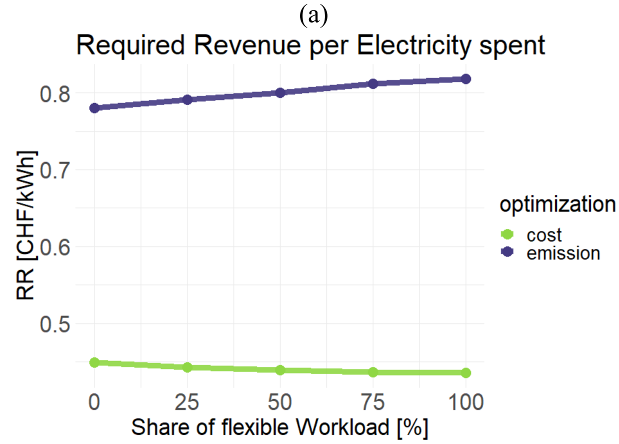

(b)

Fig. 6 KPIs values versus data center flexibility for (a) self-consumed heat fraction and (b) required revenues per electricity spent

In summary, the trend of the required revenue aligns with the optimization goal, either cost or emission and the level of flexibility in workload. While higher flexibility can lead to some cost savings for cost optimization, it may result in increased required revenue for emission optimization due to the focus on reducing carbon footprint. However, the impact of a flexible workload on the required revenue is relatively modest, and the trade-off between cost and emissions remains an essential consideration in multi-energy system.

5. CONCLUSIONS

This work analyzed the optimal integration of a small-scale data center in multi-energy building systems. From the results presented in the analysis, the following key conclusions can be derived:

- A model for the representation of the data center was successfully developed and integrated into the Ehub tool;

- Smaller data centers lead to lower required revenue per electricity spent, as electricity savings from PV installations and cost savings from exported recovered heat are more significant. However, the only waste heat recovery is not sufficient to justify the data center expense and revenues from the data center operation need to be accounted for;
- Daily flexibility in data center workload has a modest impact on self-consumed heat and required revenues. Regardless of the daily operational profile, the generated heat is fully utilized during winter, while its consumption is minimal during summer due to low or negligible domestic heating demand. The limited benefit of the IT workload flexibility can also be attributed to the availability of thermal energy storage.

ACKNOWLEDGEMENTS

The authors would like to acknowledge the financial support from the HORIZON project **ECO-QUBE**, grant agreement no. 956059, and from the Swiss Federal Office of Energy (SFOE) for the project **Sweet PATHFNDR**, grant agreement no. SI/502259-01."

REFERENCES

[1] Koot M, Wijnhoven F. Usage impact on data center electricity needs: A system dynamic forecasting model. Appl Energy 2021;291:116798. https://doi.org/10.1016/j.apenergy.2021.116798.

[2] Jinkyun C, Joonyoung Y, Changkeun L, Jinyoung L. Development of an energy evaluation and design tool for dedicated cooling systems of data centers: Sensing data center cooling energy efficiency. Energy Build 2015.

[3] Wahlroos M, Pärssinen M, Manner J, Syri S. Utilizing data center waste heat in district heating – Impacts on energy efficiency and prospects for low-temperature district heating networks. Energy 2017;140:1228–38. https://doi.org/10.1016/j.energy.2017.08.078.

[4] EcoQube 2022. https://eco-qube.eu/ (accessed June 30, 2023).

[5] Isazadeh A, Ziviani D, Claridge DE. Global trends, performance metrics, and energy reduction measures in datacom facilities. Renew Sustain Energy Rev 2023;174. https://doi.org/doi: 10.1016/j.rser.2023.113149.

[6] Ebrahimi K, Jones GF, Fleischer AS. A review of data center cooling technology, operating conditions and the corresponding low-grade waste heat recovery opportunities. Renew Sustain Energy Rev 2014;31:622–38. https://doi.org/10.1016/J.RSER.2013.12.007.

[7] Oró E, Taddeo P, Salom J. Waste heat recovery from urban air cooled data centres to increase energy efficiency of district heating networks. Sustain Cities Soc 2019;45:522–42. https://doi.org/10.1016/j.scs.2018.12.012.

[8] Ebrahimi K, Jones GF, Fleischer AS. A review of data center cooling technology, operating conditions and the corresponding low-grade waste heat recovery opportunities. Renew Sustain Energy Rev 2014;31:622–38. https://doi.org/10.1016/j.rser.2013.12.007.

[9] Andrew Bollinger L, Dorer V. The Ehub Modeling Tool: A flexible software package for district energy system optimization. Energy Procedia 2017;122:541–6. https://doi.org/10.1016/j.egypro.2017.07.402.

[10] Richner P, Heer P, Largo R, Marchesi E, Zimmermann M. NEST - A platform for the acceleration of innovation in buildings. Inf La Constr 2017;69:1–8. https://doi.org/10.3989/id.55380.

[11] BTDC One - Boden Type Data Center One n.d. https://bodentypedc.eu/ (accessed June 30, 2023).