

From servers to services: modeling data centers as heat-active urban energy prosumers


Sai Sudharshan Ravi^{*}, Theresa Sophie Löffler, Eduardo Antonio Pina, Shivom Sharma, Dorsan Lepour, Cedric Terrier, François Maréchal

Industrial Process and Energy Systems Engineering Group, École Polytechnique Fédérale de Lausanne, Rue de l'Industrie 17, 1950 Sion, Switzerland

HIGHLIGHTS

- Optimisation framework models data centers as prosumers in cities.
- Legacy vs. exergy-aware heat recovery strategies are systematically compared.
- Flexible data workloads improve PV alignment, boosting self-consumption up to 28 %.
- Data centers can meet up to 40 % of district heat demand in the case study.

GRAPHICAL ABSTRACT



ARTICLE INFO

Keywords:

Data centers
Urban energy systems
Exergy
Micro-fluidic cooling
Waste-heat recovery
MILP

ABSTRACT

The rapid expansion of cloud services has significantly increased the global energy footprint of data centers, which now account for approximately 2–3 % of global electricity consumption, a figure projected to rise to somewhere between 10 % and 51 % by 2030. While technological advances such as liquid cooling and high-temperature waste heat streams offer opportunities for improved energy efficiency, the integration of data centers into broader urban energy systems remains limited. This study investigates how data centers can transition from passive energy consumers to active prosumers through advanced heat recovery and flexible demand strategies. In this study, five system-level scenarios are modelled, varying by grid connectivity, renewable energy integration, and workload flexibility. Further, two distinct heat recovery approaches are compared: a legacy strategy that dynamically chooses between direct thermal reuse and electricity generation via an Organic Rankine Cycle (ORC), and an “exergy-aware” strategy that enforces continuous ORC operation and cascades the rejected heat from the condenser into a secondary heat pump. Using a multi-objective Mixed-Integer Linear Programming framework, the study reveals the trade-offs between data self-sufficiency and renewable energy utilization in an urban district case study comprising the EPFL campus in Lausanne. The results show that flexible computing workloads and integration with district heating networks can significantly enhance the buffering potential of data centers for variable renewable energy by up to 28 % in certain cases thereby reducing renewable curtailment, and support more efficient heat-electricity coupling by showing potential to supply up to

^{*} Corresponding author.

E-mail address: sai.ravi@epfl.ch (S.S. Ravi).

<https://doi.org/10.1016/j.apenergy.2025.127049>

Received 24 July 2025; Received in revised form 18 September 2025; Accepted 3 November 2025

Available online 7 November 2025

0306-2619/© 2025 The Authors. Published by Elsevier Ltd. This is an open access article under the CC BY license (<http://creativecommons.org/licenses/by/4.0/>).

40 % of the heat demand of the campus. This work positions data centers as critical enablers of sustainable urban energy systems and offers a transferable modeling framework for their systemic integration.

Symbols

Abbreviations	
ASHP	Air Source Heat Pump
BAU	Business As Usual
BIPV	Building-integrated Photo Voltaics
CAPEX	Capital Expenditure
CCT	Centrale de Chauffage par Thermopompe
COP	Coefficient of Performance
DC	Data Center
DHN	District Heating Network
DHW	Domestic Hot Water
ENTSO-E	European Network of Transmission System Operators for Electricity
GHG	Greenhouse Gas
GSHP	Ground Source Heat Pump
GWP	Global Warming Potential
ICT	Information and Communication Technology
MILP	Mixed-Integer Linear Programming
MST	Minimum Spanning Tree
ORC	Organic Rankine Cycle
OPEX	Operating Expenditure
PV	Photovoltaic
PUE	Power Usage Effectiveness
REHO	Renewable Energy Hub Optimizer
RCP	Research Computing Platform (context-specific cluster at EPFL)
TES	Thermal Energy Storage
TOTEX	Total Expenditure

Symbols	
$\mathcal{C}_{op}$	Operating cost (OPEX) (CHF/kWh)
$\mathcal{C}_{cap}$	Capital expenditure (CAPEX), including initial investments and replacements over the system's lifetime (CHF/kW)
$p_{s,t}^{buy}$	Retail tariff for purchasing energy carrier s from the grid at time t (CHF/kWh)
$p_{s,t}^{sell}$	Feed-in tariff for selling energy carrier s back to the grid at time t (CHF/kWh)
$E_{s,\tau}^{in}$	Energy imported from the grid for carrier s at time step τ (kWh)
$E_{s,\tau}^{out}$	Energy exported to the grid for carrier s at time step τ (kWh)
$s \in \mathcal{S}$	Set of energy carriers (e.g., electricity, gas)
$\tau \in \mathcal{T}$	Set of time steps (e.g., hourly intervals over a year)
y_u	Binary variable indicating whether unit u is installed
x_u	Installed capacity of unit u
c_u^{fixed}	Fixed cost of installing unit u (CHF)
c_u^{var}	Variable (per-kW) cost of installing unit u (CHF/kW)
c_u^{rep}	Replacement cost for unit u (CHF)
R_u	Number of times unit u must be replaced within the project horizon
r	Discount (interest) rate used for annualizing investments
N	Unit lifetime (in years)
α	Capital recovery factor for annualizing investment costs
ε	Constraint threshold applied to one of the objectives during Pareto analysis
$\mathcal{U}$	Set of all candidate energy units
ν	Frequency of occurrence of each typical/representative days of the year
GWP_{op}	Operational global warming potential (kgCO ₂ -eq.)
GWP_{const}	Global warming potential of construction (kgCO ₂ -eq.)
GWP_{tot}	Total global warming potential (kgCO ₂ -eq.)
$g_{s,\tau}^{el}$	Emission factor representing the GWP per kWh of electricity consumed at time τ (kgCO ₂ -eq./kWh of produced electricity)
g^{ng}	Emission factor representing the GWP per kWh of natural gas consumed at time τ (kgCO ₂ -eq./kWh)
$H_{s,\tau}^{in}$	Natural gas imported from the grid at time step τ (kWh)
l_u	Lifetime of a unit (years)
GWP_u^{fixed}	Fixed impact factor for GWP (kgCO ₂ -eq.)
GWP_u^{var}	Variable impact factor for GWP (kgCO ₂ -eq./kW of installed capacity)
η_{RT}	Round-trip efficiency of pumped hydro storage, assumed to be 70 %
$COP_{HP,DC}(t)$	Coefficient of performance of the ORC condenser source heat pump system at hour t
$COP_{lake}(t)$	Coefficient of performance of the lake source heat pump system at hour t

(continued on next column)

(continued)

$E_{ORC}(t)$	Electricity produced from the ORC at hour t (kWh)
$Q_{direct}(t)$	Direct heat supplied by the DC to the DHN at hour t (kWh)
$Q_{HP,DC}(t)$	Heat supplied by the heat pump linked to the data center through the condenser of the ORC (secondary heat pump) at hour t (kWh)
$\varnothing_{demand}(T, t)$	Carnot factor of the building heat demand at temperature T , time t , with reference temperature of 25 °C
$E_{HP,DC}(t)$	Electricity consumed by the heat pump linked to the data center through the condenser of the ORC (secondary heat pump) at hour t (kWh)
Q_{DC}^{total}	Total annual net heat that could be supplied by the data center (kWh)
$Q_{Demand,SH+DHW}$	Annual net heat demand of the buildings for space heating and domestic hot water (kWh)
μ	Self Sufficiency of data
$\delta_{in-house}$	Data processing that is handled locally by the on-campus data center (kWh), considered in energy equivalence
δ_{cloud}	Data processing that is outsourced, i.e., Data import (kWh), considered in energy equivalence
E_{pv}	Electricity produced by the PV installations (kWh)
E_{export}	Electricity export to the grid by the PV installations (kWh)
β	Self-consumption of the electricity produced by the PV installations

1. Introduction

The rapid advancement of cloud computing and the increasing reliance on data-intensive services have led to an unprecedented rise in the number and scale of data centers worldwide. These centers now play a pivotal role in the global economy, driving everything from e-commerce to scientific research. However, this expansion comes with significant energy costs and freshwater consumption that is required for cooling. Global server instances have climbed by 647 %, storage capacity has increased by 2500 %, and network traffic has surged by 1000 % in the past ten years [1]. While the total number of data centers is decreasing (from about 8.6 million in 2015 to 7.2 million in 2021) [2], the number of newly established hyperscale data centers (with typically 5000 or more servers that are typically used for cloud services among others) is increasing at a very quick pace. 700 of these sizable facilities were operational at the end of 2021, and an average of 16 new ones are added each quarter [3,4]. The public spending on cloud services was estimated at \$396 and \$482 billions in 2021 and 2022 respectively [5]. There is a recognizable efficiency gain from moving IT workloads to public cloud services, where computing resources, such as servers, storage, databases, and applications, are hosted by third-party providers and can be accessed over the internet. Some studies contend that this is the main reason why data center electricity consumption has plateaued over the past years, rising from 193.7 TWh in 2010 to 205.2 TWh in 2018 [1]. However, the result is still an increase in energy demand [6], from today's share of 1–2 % to up to an estimated 51 % of the world's electricity use by 2030 according to some studies [7–9]. This dramatic rise in the power consumption from data centers highlights the need for more sustainable and energy-efficient approaches to their operation, particularly in the context of their cooling systems, which often consume vast amounts of electricity and freshwater.

Unlike the electricity demands of residential and tertiary sector buildings, which fluctuate with diurnal and seasonal variations driven by occupancy, space conditioning, and lighting, data centers exhibit a near-constant energy requirement. This has prompted large-scale agreements for dedicated power supplies, such as Amazon Web Services' procurement of up to 960 MW from the Susquehanna nuclear power plant and Microsoft's 20-year power purchase agreement to revive the dormant Three Mile Island Unit 1 reactor [10]. While these agreements underscore the growing demand for reliable, emissions-free power sources, they also highlight the challenges associated with

reviving nuclear infrastructure, namely, significant capital investment, regulatory approvals, and historical safety concerns [10]. This aside, a study conducted by the Idaho National laboratory, shows that as a source of power, using small modular reactors for data centers offers a cost-effective way to power the data centers as compared to the fully-renewable counterparts due to lower projected capital costs. However, this study did not consider the effects of systemic integration or workload flexibility options with data centers [11]. Beyond electricity consumption, data centers also have a significant impact on water usage, particularly due to cooling needs. As the demand for thermal management grows with data processing, innovative cooling strategies are required, which opens new opportunities for waste heat recovery, such as the integration of data centers into district heating networks (DHNs). By harnessing thermal energy from industrial processes, waste incinerators, and existing power plants, coupled with thermal energy storage (TES) solutions, it is possible to improve the overall efficiency of urban energy systems and mitigate electricity demand for heating and cooling applications. Cooling can be provided through adsorption and absorption chillers that utilize the waste heat generated in those processes. This study proposes the application of such waste heat integration strategies to data centers, which can themselves become sources of recoverable heat. These measures can help offset grid burden and enable carbon-neutral energy pathways without the financial and operational complexities associated with reactivating nuclear facilities.

An unprecedented issue faces traditional air cooling, which has been the cornerstone of data center thermal management. It can no longer efficiently manage the rising heat generated by the ever-increasing power densities of modern IT equipment. Liquid cooling can be used to replace air as the primary cooling medium, as air-based systems become less effective and more costly at higher loads. In conventional configurations, water is often circulated near the server racks to absorb heat from the air going through them, but this process is inherently limited in its ability to recover higher-grade thermal energy. This is still not efficient as the next generation microprocessors surpass heat fluxes of 100 W/cm^2 where air can only remove about 37 W/cm^2 of the heat [12]. The decision to adopt liquid cooling is supported by a strong theoretical foundation which, regardless of the specific liquid cooling technology selected, it consistently outperforms air cooling in terms of energy efficiency. As the amount of heat that needs to be evacuated increases with a constrained area for heat transfer, a higher heat transfer rate can only be achieved with a higher temperature difference, which typically equates to higher energy spending or the use of fluids and techniques to improve the overall heat transfer coefficient of the fluid used for the heat transfer, which is where the inherent advantage of liquid-cooling technologies lie [13]. Among newer technologies, micro-fluidic chip cooling has emerged as a promising solution to improve cooling performance, while also addressing the challenge of underutilized waste heat recovery. This constitutes a paradigm shift: cooling is no longer just about keeping server temperatures to acceptable values, but also about unlocking value from waste heat, which can be reintegrated into urban energy systems. However, this transition is not without challenges, including increased infrastructure complexity, higher upfront costs, and the need for compatibility with existing systems [14,15].

In urban areas, where energy demand is steadily growing, data centers have the potential to evolve from passive energy consumers to active participants in energy production, or “prosumers” [16]. This shift requires a rethinking of waste heat management, from simple dissipation to strategic utilization. Rather than rejecting this thermal energy into the environment, it can be repurposed for heating or electricity generation, reducing the net energy consumption of both the data centers and the surrounding infrastructure.

In this study, we consider the case of the campus of École Polytechnique Fédérale de Lausanne (EPFL), located near Lake Geneva (Lac Léman) in Lausanne, Switzerland as a case study. The campus hosts its own data center and is currently planning to add a new 2 MW data

center facility adjacent to its heating plant, which supplies a district heating network powered by heat pumps [17]. This work is developed in the framework of the ‘HeatingBits’ project funded by Solutions for Sustainability Initiative (S4S) [18]. The racks in the new data center will be liquid-cooled, enabling access to higher-grade waste heat compared to traditional air-cooled systems. Therefore, we analyze the dual role of data centers as both electricity consumers and heat producers within the urban systems. The methodology followed in this study involves characterizing the campus building stock along with the data center to estimate their electrical and thermal energy demands, modeling energy generation and storage technologies capable of meeting these demands, and conducting a scenario analysis to quantify the impact of the data center on the overall energy performance of the campus energy system including thermal energy storage which plays a critical role in enabling the integration of waste heat into urban energy systems. Since the availability of waste heat from data centers may not always align with heating demand, storage becomes essential to shift thermal energy across time. In this work, we consider the district heating network (DHN) not only as a distribution system but also as a thermal energy storage buffer, enabling the campus to absorb and utilize excess heat when needed, thereby enhancing flexibility and reducing reliance on external energy sources.

The paper is structured as follows: Section 2 presents the literature review on waste heat utilization in data centers. Section 3 introduces the modeling tool and outlines the scenario definitions considered in this study. Section 4 presents and discusses the results, and finally, Section 5 provides concluding remarks, including future directions and outlook.

2. Literature review

Data centers operate continuously with a stable base load, though demand fluctuations can lead to load peaks. It must be ensured at all times that the data center’s processors are sufficiently cooled to guarantee smooth operation. According to a study of VanGeet, Lintner, and Tschudi [19] the two main methods used for cooling of the processors are air and liquid cooling. In air cooling systems cool air circulates through the data center to absorb heat from processors. In contrast, liquid cooling systems use liquids to directly absorb heat, offering higher efficiency and better suitability for high-performance computing.

Depending on the system design, the waste heat from the data center can be recovered from the exhaust air, the circulating water or the coolants used, for example [20]. Despite its potential, the relatively low temperature of this waste heat presents a challenge for effective reuse. Numerous studies have investigated strategies for recovering and utilizing this low-grade thermal energy. One of the main applications of this waste heat is heating nearby buildings, though lower temperature of waste heat remains a significant challenge, leading to further research on potential solutions.

Tervo, Syri, and Hiltunen [21] conducted a study of a 21 MW data center planned in Finland where waste heat of the data center could be recovered and used directly in the DHN. A 36 % reduction in CO₂ emissions within the district heating network (DHN) can be achieved by decreasing the reliance on biomass and peat as energy sources. Furthermore, the production cost of the district heating network could be reduced by 19 %. Another study examined the recovery of waste heat for a CO₂ Ground Source Heat Pump (GSHP) which supplies waste heat to nearby buildings. This approach incorporates a seasonal thermal energy storage strategy, where excess heat is stored in soil during non-heating periods and retrieved when needed. Compared to conventional systems, this CO₂-based GSHP demonstrates significantly improved performance, achieving up to 3.87 times higher exergy efficiency while consuming only half the electrical power [22]. A 2021 study on using trans-critical CO₂ heat pumps for recovering the waste heat from the data center’s IT room cooling water, offers a 18.2–28.9 % increase in the COP, while reducing the investment costs by up to 10.2 %. [23]

Another study compares district heating and organic Rankine cycle (ORC) systems for recovering waste heat from immersion-cooled data centers. It evaluates system performance under varying local conditions, including electricity prices, heating demand, and carbon costs. District heating is recommended in regions with high heat demand and low infrastructure costs, while ORC systems are more suitable where electricity is expensive and heating needs are low. Although chillers and heat pumps can enhance ORC efficiency, their added power consumption often reduces overall energy savings [24].

While these studies offer valuable insights into system performance, they often neglect the temporal mismatch between the steady heat output of data centers and the fluctuating heating demand of buildings. Seasonal variations further limit the usability of waste heat during warmer periods [25,26]. Although the economic feasibility of waste heat utilization has been widely studied, there is limited research on optimizing its use during non-heating periods or improving coordination between heat supply and demand. To address this gap, a 2022 study proposed a multi-mode operational strategy for combined cooling and power (CCP) systems. This approach involves switching between power generation using Organic Rankine or Kalina cycles and cooling through vapor compression, ejector expansion, or absorption refrigeration cycles, depending on seasonal and operational conditions. The study showed that a hybrid configuration combining an ORC with an absorption refrigeration cycle could achieve a maximum energy efficiency of 58.13 %, demonstrating its potential for year-round waste heat recovery. [27].

To enhance the year-round utilization of waste heat recent studies are analyzing Carnot batteries as heat recovery methods. One study proposed a liquid-cooled data center that combines a heat pump and an organic Rankine cycle. This approach enables both thermal storage and flexible energy conversion, achieving round-trip efficiencies of up to 83 % and even exceeding 100 % under favorable conditions, making it a promising solution for balancing heat recovery with fluctuating energy demands [28]. Another study explored the integration of a Carnot battery with a data center powered by a photovoltaic system, thereby enhancing the use of renewable energy. The system achieved a round-trip efficiency of 43 % and proved economically viable, with a payback period as short as 4 years under favorable electricity pricing conditions [29].

Building on these earlier approaches, several recent studies explore comparable technical pathways. Wang et al., provides design and performance guidance for high-temperature heat pumps that directly couple data center waste heat with district heating networks [30], while Corigliano et al., reviews Organic Rankine Cycle integration for digital infrastructure, focusing on fluid selection and system layouts [31]. At the optimization level, Zhu et al., demonstrates MILP-based co-dispatch of ORCs with thermal and electrical storage [32], and Chen et al., [33] formulates a multi-heat-pump recovery chain with embedded thermal storage for data centers. Complementing these, Takci et al., quantifies

the systemic value of workload flexibility as a controllable resource [34]. Together, these studies underscore a growing recognition of data centers as active energy nodes. Our work advances this line of inquiry by co-optimizing an exergy-aware ORC-HP cascade with hourly workload flexibility and district-thermal buffering in a single MILP, while mapping the CAPEX-OPEX trade-space and reporting GWP for each Pareto solution.

While the reviewed studies offer valuable insights into various approaches for recovering waste heat from data centers, a comparative perspective is essential to understand their systemic relevance and practical applicability. To this end, Table 1 provides a structured assessment of the selected studies across five key dimensions that are central to this study's objectives: systemic integration, the presence and type of storage, seasonal adaptability, the role of data centers as renewable energy buffers, and the degree of heat-electricity interconnection. These characteristics are not only indicative of the technological maturity and flexibility of each approach, but also critical for evaluating their potential contribution to resilient, low-carbon energy systems.

3. Methodology

3.1. Optimization framework

The Renewable Energy Hub Optimizer (REHO) is an open-source decision support tool designed for planning sustainable urban energy systems, developed by the Industrial Process and Energy Systems Engineering (IPESE) group from EPFL (Switzerland) [35]. REHO optimizes both the design and operation of energy capacities, considering multiple objectives such as economic, environmental, and efficiency goals. Its MILP framework adapts to various scales (building and district levels), accounts for different end-use demands, integrates multiple energy sources, and incorporates diverse energy generation and storage technologies, enabling informed decisions for renewable energy deployment in urban areas. Energy hubs or energy communities, where buildings can exchange energy and share equipment, enhance self-consumption of local resources by integrating distributed energy sources and supplying multiple services to consumers. These communities also promote energy efficiency and support high levels of renewable energy integration in urban environments, while minimizing the need for additional infrastructure and equipment upgrades.

The model considers various energy demands, including thermal comfort (space heating and cooling), domestic hot water (DHW), domestic electricity, and mobility needs. While domestic electricity, DHW, and mobility demands are pre-processed using standardized profiles based on regulations or measurements, thermal comfort demand is handled directly within the model based on hourly solar irradiation and ambient air temperature data. This approach allows the incorporation of the energy management system's control strategy and the potential for

Table 1
Comparative assessment of recent data center waste heat recovery studies across five key aspects.

Key Focus	Reference	Systemic Integration	Storage	Seasonality	DCs as Renewable Energy Buffers	Heat-Electricity Interconnection
Direct heat supply	[21]	Heating Integration	✗	✗	✗	✗
GSHP + seasonal thermal storage	[22]	Heating Integration	Thermal storage	Via storage	✗	Through heat pump
ORC + district heating	[24]	Heating and electricity integration	✗	✗	✗	ORC electricity generation
ORC + cooling modes	[27]	Heating integration	✗	✗	✗	ORC electricity generation
Carnot battery	[28]	Cooling and electricity integration	Thermal storage	Via storage	✗	Carnot battery
Carnot battery + PV	[29]	Combination of electricity and heating system	Thermal storage	Via storage	PV integration	Carnot battery
Full integration of DC waste heat	This Study	Heating and electricity integration	Thermal storage and battery storage	Via storage and ORC	PV integration	ORC electricity generation

thermal renovations of buildings. Heating and cooling needs are influenced by several factors, such as heat loss through the building envelope, the building's heat capacity, and heat gains from occupants, appliances, and solar radiation. Additionally, thermal comfort demand depends on the target room temperatures, the nominal supply and return temperatures of the heat distribution system, and its control strategy. Energy conversion technologies like heat pumps, electric heaters, fuel cells, gas boilers, and air conditioners can meet heating and cooling requirements, or it can be sourced directly from district heating systems. Energy can also be stored in devices like batteries or water tanks, or through the building's thermal inertia. Photovoltaic panels serve as renewable energy sources, while the building's energy system is connected to the district's energy distribution infrastructure, including the electrical and natural gas grids. Fig. 1 shows the different steps involved in the modeling of an urban energy hub.

The use of hourly time steps in typical representative annual profiles results in 8760 data points per year, which, when combined with the complexity of the MILP formulation, creates computational challenges that make the model difficult to track and solve. To address this, it is necessary to reduce the size of the dataset representing both the energy demand of the renewable energy hub and the corresponding weather conditions such as the hourly ambient air temperature and solar

irradiation data. One commonly used approach involves aggregating time series data into a limited number of representative periods, such as typical days or weeks, which are assigned a frequency of occurrence to reflect their distribution over the entire year.

There are many methodologies used for clustering time series data, ranging from advanced techniques such as k-means and Gaussian Mixture Models to simpler approaches based on averaging, which group time steps by chosen criteria (e.g. season, month, working/weekend day). In REHO, the k-medoids clustering algorithm is employed typically with temperature and irradiation profiles of the whole year from the specific geographical location of the urban energy system. REHO applies a Dantzig-Wolfe decomposition to solve the joint design–operation problem. The master problem represents the community level, coordinating investments and shared constraints such as transformer limits and energy exchanges, while one subproblem per building resolves the detailed hourly operation across typical periods. This decomposition approximates the full compact formulation, which would otherwise require solving all building operations simultaneously, with significantly higher computational burden. The approach reduces complexity to nearly linear growth with the number of buildings while maintaining consistency in the main objectives at both the building and system levels. Because dispatch and capacity sizing are solved in the same




Fig. 1. An overview of open-source Mixed-Integer Linear Programming tool REHO (renewable energy hub optimizer).

optimization loop, operational variability, including intra-day curtailment and flexibility requirements, directly informs technology selection. Intra-day storage resources are represented explicitly through district-level batteries on the electrical side and the thermal buffer of the district heating network in this study.

Aside from this, in this study, ICT is treated as an additional end-use demand alongside electricity, space heating, and domestic hot water. To enhance the temporal granularity of representative days, data center workloads were incorporated into the clustering process for typical day selection. For the EPFL case study, the data demand profile was aligned with that of the campus RCP cluster, which serves as a reliable proxy for overall data usage. Furthermore, spot electricity prices from ENTSO-E and the carbon intensity of grid electricity were integrated into the analysis to assess the role of renewable infrastructure in the energy system [36,37]. Therefore, in this study, typical days were identified by analyzing relevant attributes for the study such as global solar irradiation, ambient temperature, carbon intensity and electricity purchase price and the load of the data centers. For the attributes mentioned earlier and after careful consideration of different numbers of clusters, a total of 14 typical days was chosen as the most representative of the whole year. In addition, 2 extreme hours, representing severe summer and winter conditions, were included to ensure that the system configuration can handle peak stress scenarios. It should be noted that these extreme hours do not influence the system's typical operation but rather verify the system's robustness under exceptional conditions.

3.2. Energy generation and storage technologies

A multi-objective optimization problem was formulated to explore the trade-offs between operational expenses (OPEX) and capital expenses (CAPEX), for all scenarios with two heat recovery strategies. The overall global warming potential of the system configuration was computed during optimization.

$$\mathcal{C}_{op} = \sum_{s \in \mathcal{S}} \sum_{\tau \in \mathcal{T}} (p_{s,\tau}^{buy} \cdot \mathbf{E}_{s,\tau}^{in} - p_{s,\tau}^{sell} \cdot \mathbf{E}_{s,\tau}^{out}) \cdot L \quad (1)$$

$$\mathcal{C}_{cap} = \alpha \cdot \left(\sum_{u \in \mathcal{U}} (c_u^{fixed} \cdot y_u + c_u^{var} \cdot x_u) + \sum_{u \in \mathcal{U}} R_u \cdot c_u^{rep} \right) \quad (2)$$

$$\alpha = \frac{r(1+r)^N}{(1+r)^N - 1} \quad (3)$$

$$GWP_{op} = \sum_{s \in \mathcal{S}} \sum_{\tau \in \mathcal{T}} (g_{s,\tau}^{el} \cdot \mathbf{E}_{s,\tau}^{in} - g_{s,\tau}^{el} \cdot \mathbf{E}_{s,\tau}^{out} + g^{rg} \cdot \mathbf{H}_{s,\tau}^{in}) \cdot L \quad (4)$$

$$GWP_{const} = \sum_{u \in \mathcal{U}} \frac{1}{L_u} (GWP_u^{fixed} \cdot y_u + GWP_u^{var} \cdot x_u) \quad (5)$$

$$GWP_{tot} = GWP_{op} + GWP_{const} \quad (6)$$

Table 2 shows the most important technologies with their relevant characteristics utilized in modeling the system.

Table 2
Details of Important technologies utilized.

Technology	Layer in	Layer out	N [years]	c_u^{fixed} [CHF]	c_u^{var} [CHF]	GWP_u^{fixed} [kgCO ₂ eq]	GWP_u^{var} [kgCO ₂ eq]	Ref
Data Center	Electricity	Data/Heat	4	100,000	6000	0	4288	[38]
District- and building-level heat pump	Heat/Electricity	Heat	20	17,500	300	1500	15	[39]
ORC*	Heat	Electricity/Heat	20	17,500	300	1500	15	[39]
PV	–	Electricity	30	10,000	1300	200	1000	[40]
District-level battery	Electricity	Electricity	15	50,000	200	10,000	50	[41]
Electrical heater SH	Electricity	Heat	20	1000	50	100	5	[42]
Electrical heater DHW	Electricity	Heat	20	2000	50	200	5	[43]

* Proxy for ORC cost, based on component similarity with heat pumps.

3.3. Waste heat recovery strategies

To explore the potential of data centers in integrated energy systems, this study examines two distinct strategies for reusing the waste heat produced by a liquid-cooled data center. Both strategies assume a constant outlet temperature of 75 °C from DC (as has been observed from the pilot-scale demonstrations at EPFL) [44,45] and are applied consistently across all system configurations to ensure comparability. This assumption is based on the fact that the proposed data centers in the campus will use micro-fluidic cooling systems, that rise the potential heat recovery temperature from 21 °C in typical air-cooled systems to about 75 °C. This temperature is assumed to be the same across all data workloads.

The first strategy, referred to as Legacy Heat Recovery, models a hybrid system in which waste heat can be used either for space heating and domestic hot water or to drive an ORC that generates electricity. For each time step, the MILP model dynamically optimizes system operation to meet heating and electricity demands. For instance, in wintertime when thermal demand is present, the model may prioritize direct heat injection into the district heating network; while in summertime, it may divert heat to the ORC to generate electricity. This approach reflects a demand-responsive system that maximizes flexibility without enforcing year-round use of the waste heat stream.

The second strategy, Exergy-Aware Heat Recovery, enforces continuous ORC operation regardless of campus thermal demand. This ensures that waste heat is always used to generate electricity. The residual thermal energy rejected from the ORC condenser can be repurposed as a high-grade heat source for the district-level heat pump, improving its coefficient of performance. This cascading approach ensures that the full exergy content of the waste heat is exploited, achieving higher overall system efficiency at the cost of increased system complexity and infrastructure demands. Section A.2 in the Supplementary Material, visualizes both heat recovery strategies, whereas Table 3 presents main features for them.

3.4. System scenarios

Building on the heat recovery strategies, five system scenarios are considered for high levels of technological integration, operational flexibility, and energy autonomy. These scenarios vary in three key dimensions:

- Electric grid connectivity (connected or off-grid),
- District-level PV integration (presence or absence of local renewable electricity),
- Computational workload flexibility (fixed or flexible workloads). A detailed discussion of workload flexibility and the underlying assumption of this study is provided in Section 4.2.

Each scenario has two possible heat recovery strategies, resulting in total nine scenarios, as Scenario 1 has no waste heat recovery, which termed as the Business-as-Usual (BAU). This serves as the baseline configuration in which the EPFL campus uses a lake-sourced district heat pump for heating, and the data center operates independently, rejecting

Table 3

Overview of waste heat recovery strategies.

Aspect	Legacy Heat Recovery	Exergy-Aware Heat Recovery
Use of waste heat for direct heating	Allowed	Not allowed
Enhances heat pump efficiency	No. A secondary heat pump is not considered	Yes. A secondary heat pump utilizes waste heat from ORC condenser, to achieve higher COP
Strategy flexibility	Dynamic decision making on ORC operation	ORC is constrained to work at all hours of data center operation

Table 4

Summary of Scenario definitions.

Scenario	Name	Grid Connection	PV Integration	Workload Flexibility	Heat Recovery Strategy
S1	BAU	✓	✗	✗	None
S2	DC with fixed workloads, connection to grid and PV farms	✓	✓	✗	Legacy or Exergy-Aware
S3	DC with flexible workloads, connection to grid and PV farms	✓	✓	✓	Legacy or Exergy-Aware
S4	DC with fixed workloads, connection to PV farms	✗	✓	✗	Legacy or Exergy-Aware
S5	DC with flexible workloads, connection to PV farms	✗	✓	✓	Legacy or Exergy-Aware

all waste heat. Scenario 2, Grid-Connected with Fixed Workload, adds district-level photovoltaic (PV) generation to the grid-connected system while maintaining a fixed data processing schedule. Scenario 3, Grid-Connected with Flexible Workload, retains the grid connection and PV generation but allows for time-shifting of computational tasks to better align with renewable electricity availability and waste heat recovery

$$Q_{DC}^{total} = \frac{\sum_{t=1}^{8760} \left[E_{ORC}(t) - Q_{HP,DC}(t) \left(\frac{1}{COP_{lake}(t)} - \frac{1}{COP_{DC}(t)} \right) + \emptyset_{demand}(T, t) \cdot (Q_{direct}(t) + Q_{HP,DC}(t)) \right]}{\sum_{t=1}^{8760} [\emptyset_{DC}(T, t) \cdot Q_{DC}(t)]} \quad (8)$$

opportunities. Scenario 4, Off-Grid with Fixed Workload, explores full electrical autonomy by powering the data center solely with PV generation; any excess data demand is outsourced, and no flexibility is introduced in workload scheduling. Finally, Scenario 5, Off-Grid with Flexible Workload, combines off-grid operation, PV generation, flexible workload management, and waste heat recovery, representing the most integrated and autonomous system configuration designed to maximize resource efficiency. An overview of the different scenarios is given in Table 4 and a schematic of the setup in Fig. 2.

To evaluate the relative merit of different heat recovery methods

from data centers, a heat ratio was introduced as a key performance indicator. While established flexibility metrics such as ramping capability, load shiftability, and temporal mismatch indices offer valuable tools for assessing system responsiveness [46], our focus in this study is on a metric that compares the exergetic performance of the heat recovery strategies. This reflects our primary interest in evaluating thermodynamic quality rather than providing a comprehensive quantification of operational flexibility. This heat ratio defined in this study, quantifies how much of the data center's waste heat can effectively be utilized to produce useful output in terms of energy services. It is calculated as the ratio between the exergetic value of the useful outputs to the exergetic value of the heat input from the data center. The total exergetic value of the outputs, shown in the numerator Eq. (8), captures the combined contribution of different recovery pathways over all hours of the year, based on time-dependent Carnot factors calculated with reference to a temperature of 25 °C.

Specifically, it includes the electricity recovered through the Organic Rankine Cycle $E_{ORC}(t)$, which has an exergetic value of 1. The secondary heat pump that extracts condenser of the ORC to the required temperature level, while producing this heat at a higher COP still requires additional electricity. The heat produced by this heat pump is effectively replacing the heat that would otherwise be supplied by the lake source heat pump, that would be operating at a lower COP. Hence, only the relative difference in electric consumption of a heat pump using lake as the source and ORC condenser as the source was deducted from as the effective consumption of the secondary heat pump. The remaining heat terms, $Q_{HP,DC}(t)$ and $Q_{direct}(t)$ are accounted for their exergetic importance considering a time varying Carnot factor that takes into account the temperature of the supply temperature required by the buildings at time t , since this value effectively captures the useful component of the heat that is used to supply this particular energy service. Throughout, the Heat Ratio is evaluated on a net basis in which the ORC term denotes net electrical output at the unit terminals, heat terms are expressed as exergy via the time-varying Carnot factor, the secondary heat pump is debited by its incremental electricity consumption, and residual auxiliary loads are accounted for in the system power balance to ensure fair comparison and proper interpretation of system-level efficiency.

$$\text{Heat ratio} = \frac{\text{Exergetic value of the outputs}}{\text{Exergetic value of the heat input from the data center}} \quad (7)$$

3.5. Data self-sufficiency and PV self-consumption

To analyze the efficiency of integrating larger PV capacities as renewable energy supply for different sizes of the data center, two different metrics are defined.

Data Self-sufficiency: This term is introduced to estimate how much of the data processing is handled locally by the on-campus data center ($\delta_{in-house}$) and how much is outsourced (δ_{cloud}). This metric calculates




Fig. 2. A schematic view of the infrastructure available at the EPFL campus showing the possible waste heat utilization pathways from the data center.

the percentage of data processed in-house compared to the total data processed (both in-house and outsourced). It is computed using the Eq. (9):

$$\mu = \frac{\delta_{\text{in-house}}}{\delta_{\text{in-house}} + \delta_{\text{cloud}}} \quad (9)$$

PV self-consumption: This term is introduced to measure how much electricity produced by the PV is used in the data center. This metric calculates the percentage of PV energy that is used directly by the data center, excluding the energy exported to the network. It is computed using Eq. (10). Here, E_{PV} is the electricity produced by PV, whereas E_{export} is the electricity export to the grid.

$$\beta = \frac{E_{\text{PV}} - E_{\text{export}}}{E_{\text{PV}}} \quad (10)$$

3.6. Case study description: EPFL campus

To model the current infrastructure of EPFL, the current and proposed infrastructure plans on data centers at EPFL were utilized. The campus currently hosts data centers that are about 4 MW in size, with plans to extend this installed capacity by 4 MW. The campus also has shared facilities with Université de Lausanne (UNIL) located next to EPFL and has servers hosted outside of EPFL. Hence, the data workload profiles, obtained from RCP clusters, which is one of the campus data center clusters, was utilized and was scaled up to a 10 MW size. These profiles were then made to be an end-use demand that will have to be met at EPFL, by either processing them in EPFL's own data centers (in-house) or by outsourcing these data processing to cloud servers. During the course of this study, this aspect of using cloud data centers, will be referred to as data import to maintain a coherence in terminology across the different resources considered in the study. While the demand

profile considered might be an overestimation of the actual data demand of EPFL, this was done to err in the side of caution, as ICT demand of EPFL, with the growth of large language models, is expected to grow and this demand profile would still be a valid assumption in this case. In modeling the data center of EPFL, a PUE of 1.3 was assumed. PUE is a metric that measures the energy efficiency of a data center, defined as the ratio of total facility energy to the energy used by IT equipment. This defines the amount of electricity needed to process one unit of data within the campus data center. Almost 95 % of the electricity consumed by the data center is assumed to be dissipated as heat that is recoverable with 5 % assumed to be lost to the cooling system of the server rooms hosting the data center [47]. The campus has an extensive DHN characterized by two different heat supply levels depending on the buildings'

Table 5

Campus energy and data demands across the typical days considered.

Typical day d	Frequency of typical days of Type d per year	Space Heating demand (MWh/day)	Domestic Hot water demand (MWh/day)	ICT demand (MWh/day)
1	29	132.41	2.15	69.68
2	28	165.71	2.61	61.96
3	27	203.44	2.46	69.68
4	18	172.38	2.46	72.51
5	23	147.07	2.15	58.84
6	4	145.88	2.46	75.64
7	29	131.55	1.16	66.35
8	39	76.01	2.76	51.31
9	18	102.59	1.05	70.71
10	39	29.34	2.76	65.90
11	24	0	2.16	70.91
12	38	0	1.05	61.96
13	25	27.87	2.46	61.68
14	24	0	1.25	61.63

construction period. Newer buildings require heat at a lower temperature (at about 40–45 °C), while older buildings require higher temperature supply (at about 65–70 °C). However, in this study, it was assumed to be a single network for simplicity and has been discussed in detail in Section 4.1. This DHN is connected to the heat pumps of the campus and will be used to supply both the space heating and domestic hot water demands of the campus buildings. Table 5 shows the summary of the space heating, domestic hot water and ICT demands of the campus across the 14 representative days and the frequency in which they appear in a given year.

4. Results and discussions

4.1. Role of the DHN as intra-day thermal storage buffers

Complementarily to the waste heat recovery strategies proposed in Section 3.2, we propose to explore the buffer potential for TES within the DHN itself. This operation can benefit, on the one hand, from the direct waste heat recovery from the DC, and on the other hand, from dynamic heat pump operation strategies.

As a first approach, it was decided to model a single DHN for the entire EPFL campus with a maximum supply temperature of 70 °C. This temperature was selected, first, to ensure compliance with public health guidelines, which recommend storing water between 60 and 70 °C for thermal disinfection (for example, from *Legionella* bacteria), and second, because it enables lower distribution losses, ensuring sufficient temperature levels at the end-users' point of use, while also avoiding oversized infrastructure. Given the geographical distribution of EPFL's buildings, data center and heat pump across the campus, as well as the buildings' heating demands, the optimal network layout was estimated using the Minimum Spanning Network (MSP) algorithm [48]. This allowed for a precise estimation of the direction of the pipes, along with their lengths and diameters. An open street map view of the campus can be seen in Fig. 3(a), where the building CCT marks the location of the data center and the heat pumps. The line thickness correlates with pipe diameter, which was estimated for the said supply temperature and heating demands as explained in the supplementary material section. From the lengths and diameters of the different pipes in the DHN, the overall volume of water stored in the network was estimated to be about

781 m³ offering a storage buffer potential of up to 18.2 MWh within a temperature difference of 20 °C. This information was then used to assess the role of DHN as thermal buffers with a simple formulation of the charged and discharged energy, as well as thermal losses, within the pipeline as depicted in Fig. 3(b). In this way, charging is represented by an increase in the temperature of the water inside the pipes. It should be noted that representing the DHN as a zero-dimensional (0-D) thermal buffer is appropriate for the planning scale of our MILP, where the primary decisions are technology selection and hourly co-dispatch (PV, ORC and HP cascades, battery, workload shifting, DHN buffering). Fine-grained network physics such as transport delays, mixing/stratification, and distributed heat losses mainly affect hydraulic operation and sub-hourly transients, which are beyond the scope of this study. Although this abstraction can modestly misrepresent rapid temperature propagation or short-cycling constraints, it preserves the first-order energy balance that determines the trade-offs analyzed here. As a natural extension, the planning model could be coupled with a transient DHN solver to assess detailed operational behavior for selected Pareto solutions.

The simulation results illustrate how the MILP model effectively leverages the DHN as an intra-day thermal storage system in response to dynamic electricity prices. Fig. 4 shows this trend for typical days in the winter and summer. The translucent red bands show the regions increase in electricity prices and a corresponding discharging trend while the translucent green bands show the opposite trend. In wintertime, the DHN is charged predominantly during low-tariff hours in the early morning (0–6 h), with stored heat discharged around hour 9, coinciding with peak electricity prices. Similar charging behavior can also be noticed between hours 10–14, when the tariffs are at a local minimum. This behavior suggests the model prioritizes avoiding high-cost electricity imports or expensive heat generation technologies during peak hours, thereby minimizing operational expenditure. By contrast, in summertime, a more complex pattern emerges due to the presence of two daily price peaks. The DHN discharges during the morning price increase (5–7 h) and recharges during the mid-day price trough (11–13 h), reflecting cost-optimal operation. The evening storage reduction corresponds to the second price peak, indicating strategic discharging.

Overall, the model demonstrates consistent behavior, adapting to seasonal electricity price signals by optimizing thermal energy storage




Fig. 3. (a) Estimated DHN layout of EPFL campus; (b) TES buffer modeling.




Fig. 4. Intra-day thermal storage buffer in the DHN characterized by charging and discharging profiles with a proportionate increase or decrease in the DHN temperature respectively in line with electricity tariffs in a given day.

operation to flatten cost profiles and enhance system efficiency. The response to dynamic electricity prices, which characterizes this optimal operation strategy, highlights the need for effective forecasting of not only day-ahead market prices of grid electricity, but also day-ahead weather conditions, which directly influence the buildings' heating demand. This possibility of using the DHNs as a buffer was then allowed for the subsequent scenarios discussed in this study.

4.2. Workload flexibility in data centers

The workloads used in this study were derived from real-time operations of the RCP data center clusters at EPFL, providing a good approximation for the campus ICT demands. These power profiles reflect energy consumption associated with processing specific data instances over defined time intervals. However, not all data processing is time critical. With the growing adoption of large language models




Fig. 5. Original and shifted workload profiles of the data center alongside the corresponding solar irradiation profiles for a given week of the year. Each profile is normalized to its maximum value over the entire year to allow for direct comparison.




Fig. 6. Pareto results for the “Business as usual” scenario S1, plotted as a function of in-house data processing.

(LLMs), it has become increasingly common for data centers to shift non-urgent workloads to periods of lower server traffic. The extent to which workloads can be shifted depends on the nature of the data being processed. Due to limited access to detailed information on workload composition in data centers worldwide, values from the literature were used to give a range of shiftable load, from 30 % to 70 % of the total workload at any given time interval [34,49,50].

To evaluate the potential benefits of flexible workloads, a conservative estimate of 30 % shiftable load was assumed, with a shifting horizon of 24 h covering the entire representative day. This implies that while the total daily workload remains constant, its distribution within the day can be adjusted to better align with PV production and time-varying electricity purchase prices. To model this, an objective function was defined to minimize the sum of squared hourly differences between the normalized shifted data center load and the normalized solar irradiance profile:

$$\text{Minimize } \sigma = \sum_{t=1}^T (S_t - I_t)^2 \quad (11)$$

where:

- T is the total number of time steps (e.g., hours in a year),
- S_t is the normalized shifted load at time t ,
- I_t is the normalized irradiance at time t .

A flowchart outlining the process of load shifting in line with PV generation can be found in the Supplementary Material. Fig. 5 shows an example of a representative week where the workload has been shifted based on the irradiation profile.

4.3. Scenario analysis

Multi-objective optimization (MOO) analyses were conducted for

each of the five scenarios to explore trade-offs between economic objectives. In REHO, the ϵ -constraint method constructs the Pareto front by selecting one objective as the primary target and constraining the other with progressively tighter bounds. More explanation on the methodology can be found in [51,52]. In our case, CAPEX was chosen as the primary optimization objective, while OPEX was progressively bounded by ϵ values linearly distributed between upper and lower limits. The upper limit is the OPEX value obtained for the minimum CAPEX, while the lower limit corresponds to the optimal OPEX. The number of intermediate steps was controlled by the parameter that selects the number of Pareto points chosen during the optimization, yielding a set of non-dominated solutions spanning the entire frontier. For each solution, the associated Global Warming Potential (GWP) was also calculated, enabling a three-dimensional assessment of trade-offs. The representative solutions discussed in the results were manually selected from this set for discussion in the subsequent sections.

To ensure consistent comparison across scenarios, all results are normalized by the building's energy reference area (ERA) and expressed in CHF or kg CO₂ per m² per year. This allows data, heat, and electricity demands to be evaluated on a common basis, independent of building size or configuration. Detailed results for the off-grid scenarios are provided in the supplementary material to maintain clarity in the main discussion. It must also be noted that the metrics selected for comparison in this study were primarily chosen to align with the scope and structure of the urban energy system model. A comparison of economic indicators has not been carried out, as this would require the definition of a consistent economic framework across all system scenarios or a pre-selection of certain scenarios for this reason. An economic assessment has hence been deemed to be outside the methodological focus of this work. Instead, the study prioritizes technical and systemic performance metrics discussed before to draw robust and comparable insights.




Fig. 7. Composite curves for typical winter (a, b) and summer (c, d) hours comparing direct (a, c) and exergy-aware (b, d) heat recovery strategies. Curves are plotted using Carnot factors to visualize exergy destruction (red area). The exergy-aware strategy, involving ORC and heat pumps, shows improved thermal matching and reduced exergy destruction compared to direct heat use, particularly in summer. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

4.3.1. S1: Business-as-usual (BAU)

The business-as-usual scenario, attempts at meeting the demands of the system on the electricity and thermal energy sides, without any form of meaningful heat recovery, as is the case today. The multi-objective optimization first minimized CAPEX and OPEX independently, thereby identifying the extreme points that determine the bounds of the Pareto front. Subsequently, a series of optimizations were performed in which CAPEX was minimized while progressively imposing an upper bound (ε-constraint) on OPEX.

Fig. 6 illustrates the results of the multi-objective optimization. The solution at 0 MJ_{el}/m²_{ERA} in-house data processing (left-most point) corresponds to the minimum CAPEX, and the solution at 250 MJ_{el}/m²_{ERA} in-house data processing corresponds to minimum OPEX. Electricity import, electricity export, and data imports are associated with OPEX, while investments in data center, PV panels, heat pump, and district-level battery are associated with CAPEX. CAPEX minimization results in a huge import of both electricity and data, while OPEX minimization results in installation of maximum size of data center that is needed to in-house process all data of EPFL. In the baseline scenario (without a data center) operational costs are highest, with around 10 CHF/m²/year for electricity and 30 CHF/m²/year for data, as all services are imported.

Installing a data center initially leads to a noticeable reduction in total costs by lowering data import needs and improving energy efficiency. The minimum TOTEX, representing the sum of CAPEX and OPEX, was reached around 245 MJ_{el}/m²_{ERA} of in-house data processing, indicating the benefit of allowing for a very small flexibility margin between data outsourcing and in-house processing. Fig. 6 also presents the GWP emissions for the different Pareto solutions, where high OPEX leads to larger emission values. Hence, GWP for the overall system progressively reduces with the reduction in the amount of imported data. It is interesting to note that the minimum GWP is reached for the same minimum TOTEX point. (See Fig. 6.)

4.3.2. Exergy destruction across the heat recovery strategies

Before analyzing the full-year performance of the two heat recovery strategies, a targeted heat integration assessment was conducted for representative hours during a typical winter and summer day. Composite curves were constructed for each strategy to evaluate thermal match quality and exergy efficiency as shown in Fig. 7. In the winter scenario, the curves reveal a significant temperature mismatch when waste heat from the data center is directly transferred to the district heating network (DHN), indicating limited thermodynamic




Fig. 8. Pareto results for Scenario S2. DC with fixed workloads, connection to grid and PV farms, and (a) legacy heat recovery and (b) exergy-aware heat recovery.




Fig. 9. Pareto results for Scenario S3. DC with flexible workloads, connection to grid and PV farms, and heat recovery.

compatibility. In contrast, the exergy-aware strategy which first converts the waste heat to electricity via an Organic Rankine Cycle (ORC) and subsequently uses this electricity to drive heat pumps exhibits a smaller temperature gap, signifying improved thermal integration. In this configuration, building heat supply is met through a combination of a lake-sourced primary heat pump and a secondary heat pump that recovers low-grade heat from the ORC condenser. Importantly, the area

between the hot and cold composite curves, plotted against Carnot factors rather than absolute temperatures, serves as a visual and quantitative indicator of exergy destruction. This area is notably larger in the direct use strategy, particularly during summer, when heat demand occurs at very low temperatures, thereby exacerbating exergy losses. These insights underscore the thermodynamic advantages of the exergy-aware strategy, especially under conditions of low-temperature demand.

4.3.3. S2: DC with fixed workloads, connection to grid and PV farms

In this scenario, district-level PV panels are modelled to enhance the supply of renewable electricity. A summary of the scenarios' assumptions was given in Table 4. As illustrated in the Pareto front between CAPEX and OPEX in Fig. 8, the highest total annual costs are observed in the case without a data center, reaching 50 CHF/m²/year. In this case, data import dominates the cost breakdown, contributing approximately 35 CHF/m²/year. As the data center size increases, the reliance on external data import decreases, leading to a significant reduction in both capital and operational expenditures. In fact, increasing the in-house processing of data leads to a noticeable increase in the size of the district-level PV systems, a decrease in the imported electricity, and an increase in the amount of electricity exported to the grid. This is because of the data workload demand profiles that are typically not synchronized with the PV availability.

The optimal total annual cost is observed again high in-house data processing of up to 250 MJ_{el}/m²_{ERA}. Interestingly, the GWP is lowest at the same point, falling below 50 kg CO₂/m²/year, highlighting a strong interest of this solution in terms of both financial and environmental performances. It can also be noted that the overall GWP of the system starts to climb towards higher sizes of data center, as the PV installation reaches its maximum installable size. Additionally, the heat ratio for the exergy-aware heat recovery strategy can be seen to reach around 0.6, indicating a relatively effective use of heat rejected by the data center as opposed to a heat ratio of 0.4 in the legacy heat recovery strategy. In general, the heat ratio increases with increasing amount of heat recovery in terms of electricity generation, since this will have the highest exergetic value. The heat recovery ratio in the exergy heat recovery strategy appears to be relatively constant, since ratio of the exergetic value of the output generated to the input is determined by the configuration of the system, which is fixed a-priori in this strategy.

4.3.4. S3: DC with flexible workloads, connection to grid and PV farms

In scenario S3, the shifted workloads were implemented in REHO, and the corresponding results are shown in Fig. 9. As the system tries to meet the shifted ICT demand with PV generation, it significantly over-sizes the data center, reaching up to twice the size compared to previous fixed-demand scenarios. This means that as the emphasis on reduction of operation expenses increases, an increase in renewable integration is preferred which leads to configurations that require data centers that are bigger than the base size of 10 MW that is considered across this study. The electrical grid continues to play a crucial role, enabling in-house data processing during hours of low or no PV generation, resulting in very low levels of data imports.

The GWP profile of the overall system in both heat recovery strategies initially shows a progressive decline. However, as the size of the data center continues to increase, the embodied emissions from its construction begin to have a more noticeable impact on the system's total GWP. Another important observation related to electricity exports, is that when considering the similar data center capacities configurations between the fixed and the flexible workloads, despite an increased installed capacity of the PV in the flexible workload scenario, a higher self-consumption of the generated renewable electricity and reduced curtailment, improving overall energy utilization is observed.

The optimal total annual cost, defined as the sum of CAPEX and OPEX, is achieved at 160 MJ_{el}/m²_{ERA} of in-house data processing. Interestingly, this point also corresponds to the lowest GWP, suggesting a strong alignment between economic and environmental objectives. While heat ratio trends across the two strategies remain similar to that as in Scenario S2, it must be noted that the costs of the data center have increased significantly in this scenario due to the oversizing to meet the peaks of data demand.

The optimization results also reveal a distinct cost pattern: A higher




Fig. 10. Sankey diagrams showing the system configurations across the two heat recovery strategies in the grid-connected scenarios S2 (a, b) and S3 (c, d).

total system costs occur at both ends of the in-house data processing range, corresponding to minimum OPEX and minimum CAPEX. Initially, costs are driven by high levels of data and electricity imports. At the upper end, costs rise due to the extensive deployment of technologies, including PV panels, district-level PV systems, district batteries, heat pumps, and the data center itself. At about $160 \text{ MJ}_{\text{el}}/\text{m}_{\text{ERA}}^2$ in-house data processing, a sharp increase in the deployment of these technologies is observed, leading to significantly higher system costs.

4.4. Energy flow analysis for scenarios S2 and S3

Analyzing energy flows is important to understand how energy is converted, distributed and recovered within the system. It helps to identify inefficiencies, quantify the impact of different operating strategies and assess the effectiveness of energy recovery measures. In the context of scenarios S2 and S3, such an analysis provides valuable insights into how flexible workloads and heat recovery strategies affect overall system performance and energy consumption.

The Sankey diagrams shown in Fig. 10(a)–(d), illustrate the energy flows for Scenarios S2 and S3. These scenarios have connections to grid, installed PV farms, two heat recovery strategies, and fixed (Scenario S2) and flexible (Scenario S3) workloads. In the legacy heat recovery strategy, the excess heat from the data center is supplied directly to the DHN and ORC. In contrast, the exergy-aware heat recovery strategy uses

excess heat solely for the ORC, with the ORC condenser supplying heat to the DHN heat pump. These diagrams are presented for 80 % of the total data processing needed to be in-house at $200 \text{ MJ}_{\text{el}}/\text{m}_{\text{ERA}}^2$, which roughly represent the current operational strategies of the campus of EPFL.

Two main observations emerge from the exergy-aware heat recovery strategy across both scenarios:

- It enables the ORC to generate approximately three times more electricity than in the legacy heat recovery strategy, as all the excess heat from the data center is directed to the ORC.
- It supplies more heat to the DHN using the data center's heat pump, whereas the legacy strategy delivers heat directly from the data center to the DHN, resulting in a lower overall heat supply contribution of the data centers to the buildings.

As a result, the legacy strategy relies more heavily on district-level heat pumps to meet heating demand. It is important to note that while the exergy-aware strategy uses the data center's heat pump, which requires electricity, the legacy strategy does not require additional electricity for heat delivery from the data center. Nevertheless, the exergy-aware strategy is more thermodynamically efficient, as it better aligns heat use with its exergy content, maximizing the value extracted from the available waste heat.




Fig. 11. Data self-sufficiency and PV self-consumptions in scenarios S2 (a), S3 (b), S4 (c) and S5 (d).

When comparing the fixed and flexible workload scenarios, the data center in the shifted case shows a lower annual electricity demand (approximately 1500 MWh less). This reduction is due to a portion of the workload being outsourced to cloud-based servers, which reduces local electricity consumption. Additionally, electricity exports are lower despite increased PV generation, due to higher self-consumption: the shifted demand profile allows for greater direct use of PV electricity instead of importing from or exporting to the grid. The alignment between demand and PV production also reduces battery charging needs, as more energy can be used in real time, further enhancing system efficiency. This behavior exemplifies the “information battery” concept, in which data centers act as buffers for renewable energy, storing excess generation in the form of processed data instead of physical energy [53].

However, in the shifted workload scenario, the data center also supplies less heat, both to the DHN and to the ORC for power generation. This reduction is due to the lower electricity demand and reduced operating hours of the data center, resulting in decreased heat output. Consequently, a greater share of the heat demand (around 4000 MWh annually) is met by heat pumps at the building level. The final heating demand remains largely unchanged, with only minor variations attributed to differences in typical day profiles derived from the clustering process.

4.5. Role of flexible workloads on data self-sufficiency and PV self-consumption

To evaluate the impact of flexible workloads on system performance, the share of data processed locally (Data self-sufficiency) and the proportion of PV electricity used directly on-site (PV self-consumption) are analyzed. These values illustrate how workload flexibility can reduce reliance on external cloud services and increase the utilization of locally generated solar energy. Fig. 11 presents the PV and data center capacities, as well as PV self-consumption, across different levels of data self-sufficiency for Scenarios S2, S3, S4 and S5.

General trends in the Pareto fronts of Scenarios S4 and S5 show that, due to the absence of a grid connection, the data centers operate solely on PV electricity, which limits their operational hours. As a result, a significant share of data must be outsourced to cloud services, leading to higher GHG emissions and relatively high system costs, mainly due to expensive data imports and reduced heat recovery potential. These dynamics are illustrated in more detail in the supplementary figures provided in the appendix.

Fig. 11 (a) illustrates the grid-connected scenario with fixed workloads S2. As the size of the data center increases, the installed PV capacity grows nearly linearly. However, due to the temporal mismatch between the fixed load and the PV generation profiles, PV self-consumption is limited to around 70 %, resulting in higher electricity exports to the grid. Data center size increases with data self-sufficiency and can reach up to 100 % data self-sufficiency, as the grid ensures continuous electricity availability. In Fig. 11 (c), the off-grid scenario with fixed workloads S4 is presented. Here, a similarly linear increase in PV capacity is observed. PV self-consumption reaches up to 90 %, as the absence of a grid connection limits the possibility of exporting excess electricity. Nevertheless, data self-sufficiency remains constrained and peaks at approximately 45 %, due to the lack of external power supply.

Fig. 11(b) shows the grid-connected scenario with flexible workloads S3. In this case, load shifting significantly improves the temporal alignment between PV generation and demand, leading to higher PV self-consumption, up to 90 %, and a more efficient use of the installed PV capacity. The installed capacity is considerably higher than in the fixed workload case, as more solar energy can be utilized effectively. Data center size increases with data self-sufficiency and reaches 100 % data self-sufficiency, supported by the grid as a reliable energy source as was observed in the other grid-connected scenario. Fig. 11(d) depicts the off-grid scenario with flexible workloads S5, which exhibits the highest PV self-consumption and the largest increase in installed PV capacity -

approximately five to ten times more than in fixed-load scenarios. The flexibility in workload scheduling allows for better utilization of solar peaks and reduces curtailment. As a result, data self-sufficiency increases significantly and reaches up to 60 %, the highest among off-grid configurations.

Across all scenarios, increasing the size of the data center leads to higher PV self-consumption and greater data self-sufficiency. Flexible workload strategies consistently outperform fixed ones, particularly in off-grid settings, by enabling more efficient use of locally generated solar energy. While grid-connected scenarios can achieve full data self-sufficiency due to the availability of external electricity, off-grid scenarios benefit substantially from flexibility, which raises self-sufficiency from 45 % to 60 %. These findings underscore the critical role of workload flexibility in enhancing energy autonomy, reducing reliance on external cloud services, and maximizing the integration of renewable energy in data center operations.

Understanding the relationship between installed PV panels and data center capacity is critical to optimizing energy efficiency and sustainability, as it helps determine how much capacity is needed to power data centers. To determine how much PV capacity is needed per unit of data center capacity, the installed PV capacity is multiplied by the self-consumption rate and divided by the data center capacity. In the fixed workload scenario, 8–9 units of PV per unit of data center capacity are required. In the flexible workload scenario, this ratio decreases from 7.8 to 5.2 as the data center size increases. This shows that flexibility reduces the PV capacity needed to meet energy demands, improving overall system efficiency.

5. Conclusions and future outlook

This study analyzed several scenarios to assess the role of data centers as energy prosumers in urban energy districts. Two heat recovery strategies were compared: (i) a direct deployment of waste heat into the DHN of EPFL campus, and (ii) a novel exergy-aware approach, which involves an ORC to convert data center waste heat into electricity and a secondary heat pump to recover rejected heat from the ORC condenser benefitting from a higher COP. The potential of data centers as renewable energy buffers, by exploiting data demand flexibility to mitigate PV curtailment, was also explored. Furthermore, the DHN was examined for its role as a thermal energy storage buffer, contributing to a holistic understanding of the energy system. Some of the key findings of this study are summarized as follows:

1. The heat ratio, introduced as a KPI to compare heat recovery strategies shows that the heat recovery ratio in the exergy heat recovery strategy appears to be relatively constant due to a fixed system configuration and shows up to 60 % of heat rejected from the data centers can be used to meet an energy service.
2. The introduction of soft and hard constraints on district-level PV electricity use revealed that data centers can act as buffers for intermittent renewables under both flexible and fixed workload profiles improving self-consumption of PV up to 90 % in the flexible cases. Interestingly, the model favored oversizing the data center to reduce operational expenditure, especially in scenarios with flexible loads. This was consistent across both grid-connected and off-grid configurations.
3. In off-grid scenarios with flexible workloads, a portion of the data demand must still be outsourced to cloud servers, despite the oversized local data center which only achieves up to 60 % data self-sufficiency. While seemingly less attractive, this configuration supports sustainability goals and could become viable with geographic coordination and high renewable availability. This strategy enables greater PV penetration and greener data processing at the cost of higher capital investment.
4. The DHN demonstrated an untapped potential as a short-term thermal buffer by allowing a buffer potential of 18.2 MWh. By

connecting more data centers to the DHN, it is possible to optimize system operation for multiple objectives, including total annual cost and GWP, while enabling flexible data center operation that aligns with peak heating demands.

5. The analysis of PV integration and workload flexibility highlighted the potential of data centers as active prosumers. Larger data centers exhibited higher PV self-consumption and data self-sufficiency, especially when paired with flexible workloads. In off-grid scenarios, data self-sufficiency increased from 45 % (fixed) to 60 % (flexible workloads), while PV self-consumption remained high in both cases (up to 90 %). Moreover, flexible workloads also reduced the required PV capacity per unit of data center capacity from around 8–9 units in fixed workload scenarios to 5.2 units, indicating more efficient use of renewable resources.

This study applied simplifying assumptions for workload flexibility, as the nature of data instances and their associated energy consumption is typically confidential for privacy and security reasons.

5.1. Policy implications

This study assessed the systemic role of data centers as thermal and digital prosumers within urban energy systems. By modeling multiple heat recovery configurations, ranging from no heat recovery to exergy-aware ORC-HP systems, and their integration into a university campus district, we demonstrated the dual potential of data centers to: (1) deliver dispatchable low-carbon heat, and (2) act as buffers for surplus renewable electricity through flexible ICT loads. Key insights are as follows:

- Exergy-aware heat recovery (e.g., ORC + HP) achieves up to 60 % higher heat recovery rates and significant GHG savings compared to legacy cooling systems, particularly in solar-intensive regions with DHNs. These benefits are increased with lower supply temperature requirements, highlighting the relevance of low-temperature DHNs in new or renovated urban energy systems.
- ICT load flexibility enables better alignment with PV production and cooling needs, improving self-sufficiency and reducing curtailment.
- Emissions from data outsourcing are non-negligible. We assumed an emission factor of 1.6 kgCO₂/kWh of imported data, underscoring the need for transparent accounting of digital infrastructure's carbon footprint.

These findings have clear implications for urban energy planning:

1. **Urban integration:** Treat data centers as active nodes in energy planning, especially in cities with district heating or ambitions for zero-carbon infrastructure. Urban masterplans should exploit thermal synergies between data centers and nearby buildings to boost energy and land-use efficiency.
2. **Policy coordination:** Require transparency from data center operators, regarding their infrastructure, energy sourcing, and workload distribution, to enable accurate and consistent tracking and quantification of GHG emissions, particularly in Scope 3 emissions accounting for digital services. The information can be used to guide more informed planning, incentivize local data processing or tax high-emission cloud imports to internalize environmental costs.
3. **Zoning and permitting:** Encourage co-location of data centers near thermal loads (e.g. urban districts) and promote designs that allow for modular thermal coupling and reversible cooling-heating systems.
4. **Flexibility markets:** Integrate digital flexibility into grid balancing and demand-response strategies. This requires effective coordination between Transmission System Operators (TSOs), Distribution System Operators (DSOs), data center operators, and renewable infrastructure managers.

CRedit authorship contribution statement

Sai Sudharshan Ravi: Writing – original draft, Visualization, Software, Methodology, Formal analysis, Conceptualization. **Theresa Sophie Löffler:** Writing – review & editing, Visualization, Software. **Eduardo Antonio Pina:** Writing – review & editing, Project administration. **Shivom Sharma:** Writing – review & editing, Investigation. **Dorsan Lepour:** Writing – review & editing. **Cedric Terrier:** Visualization. **François Maréchal:** Supervision.

Funding

This research is carried out in the frame of the “HeatingBits: renewable-supplied data centers integrating heating and cooling supply of local districts” project with the financial support of the EPFL Solutions4Sustainability program.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Acknowledgement

During the preparation of this work the authors used ChatGPT to improve the language used in this tool. After using this tool/service, the authors reviewed and edited the content as needed and takes full responsibility for the content of the publication.

Appendix A. Supplementary data

Supplementary data to this article can be found online at <https://doi.org/10.1016/j.apenergy.2025.127049>.

Data availability

Data will be made available on request.

References

- [1] Masanet E, Shehabi A, Lei N, Smith S, Koomey J. Recalibrating global data center energy-use estimates. *Science* 2020;367(6481):984–6. <https://doi.org/10.1126/science.aba3758>.
- [2] AI software market to quadruple to \$251B by 2027 – Computerworld. Accessed: Sep. 21, 2024. [Online]. Available: <https://www.computerworld.com/article/1638387/ai-software-market-to-quadruple-to-251b-by-2027.html>.
- [3] Hyperscale data center capacity doubles in under four years; the US still accounts for half | synergy research group. Accessed 2024 [Online]. Available: <https://www.srgresearch.com/articles/as-hyperscale-data-center-capacity-doubles-in-under-four-years-the-us-still-accounts-for-half-of-the-total>.
- [4] Hyperscale data center count grows to 659 – ByteDance joins the leading group | synergy research group. Accessed 2024 [Online]. Available: <https://www.srgresearch.com/articles/hyperscale-data-center-count-grows-to-659-bytedance-joins-the-leading-group>.
- [5] Gartner says four trends are shaping the future of public cloud. Accessed 2024 [Online]. Available: <https://www.gartner.com/en/newsroom/press-releases/2021-08-02-gartner-says-four-trends-are-shaping-the-future-of-public-cloud>.
- [6] 2021 Review – Microsoft and Amazon Dominate IT Vendor Revenue and Growth Ranking | Synergy Research Group. Accessed: Sep. 21, 2024. [Online]. Available: <https://www.srgresearch.com/articles/2021-review-microsoft-and-amazon-dominate-it-vendor-revenue-and-growth-ranking>.
- [7] Ferreboeuf H, et al. Lean ICT: towards digital sobriety. In: *Report for the Think Tank The Shift Project*. Vol. 6; 2019. p. 16–28.
- [8] Obringer R, Rachunok B, Maia-Silva D, Arbabzadeh M, Nateghi R, Madani K. The overlooked environmental footprint of increasing internet use. *Resour Conserv Recycl* 2021;167:105389. <https://doi.org/10.1016/j.resconrec.2020.105389>.
- [9] Andrae A, Edler T. On global electricity usage of communication technology: trends to 2030. *Challenges* 2015;6(1):117–57. doi:10/gftvrb.
- [10] Data center owners turn to nuclear as potential electricity source - U.S. Energy Information Administration (EIA). Accessed: Nov. 15, 2024. [Online]. Available: <https://www.eia.gov/todayinenergy/detail.php?id=63304>.

- [11] Soto Gonzalez GJ, et al. Powering data centers with clean energy: a techno-economic case study of nuclear and renewable energy dependability. Idaho Falls, ID (United States): Idaho National Laboratory (INL); 2024. INL/RPT-24-79307-Rev000, Aug. Accessed: Jul. 09, 2025. [Online]. Available: <https://www.osti.gov/biblio/2438321>.
- [12] Habibi Khalaj A, Halmagum SK. A review on efficient thermal management of air- and liquid-cooled data centers: from chip to the cooling system. *Appl Energy* 2017; 205:1165–88. <https://doi.org/10.1016/j.apenergy.2017.08.037>.
- [13] Huang D, Yang C, Li B. Rack-level cooling and server-level cooling. In: Data center handbook: plan, design, build, and operations of a smart data center; 2021. p. 227–37. <https://doi.org/10.1002/9781119597537.ch14>.
- [14] Chen H, Li D. Current status and challenges for liquid-cooled data centers. *Front Energy Res* 2022;10:952680.
- [15] Azarifar M, Arik M, Chang J-Y. Liquid cooling of data centers: a necessity facing challenges. *Appl Therm Eng* 2024;247:123112.
- [16] Huang P, et al. A review of data centers as prosumers in district energy systems: renewable energy integration and waste heat reuse for district heating. *Appl Energy* 2020;258:114109. <https://doi.org/10.1016/j.apenergy.2019.114109>.
- [17] Thermal power plant. EPFL. Accessed: Sep. 23, 2024. [Online]. Available: <https://www.epfl.ch/about/sustainability/energy-buildings/thermal-power-plant/>.
- [18] HeatingBits Project. EPFL. Accessed: May 25, 2025. [Online]. Available: <https://www.epfl.ch/research/solutions-for-sustainability-initiative-s4s/heating-bits/>.
- [19] Van Geet O, Sickinger D. Best Practices Guide for Energy-Efficient Data Center Design. National Renewable Energy Laboratory (NREL); Jul. 2024. Accessed: May 27, 2025. [Online]. Available: <https://www.osti.gov/servlets/purl/2417618/>.
- [20] Yuan X, Liang Y, Hu X, Xu Y, Chen Y, Kosonen R. Waste heat recoveries in data centers: a review. *Renew Sust Energ Rev* 2023;188:113777. <https://doi.org/10.1016/j.rser.2023.113777>.
- [21] Tervo S, Syri S, Hiltunen P. Reducing district heating carbon dioxide emissions with data center waste heat – region perspective. *Renew Sust Energ Rev* 2025;208:114992. <https://doi.org/10.1016/j.rser.2024.114992>.
- [22] Wang X, et al. Energy, exergy, and economic analysis of a data center energy system driven by the CO₂ ground source heat pump: prosumer perspective. *Energy Convers Manag* 2021;232:113877. <https://doi.org/10.1016/j.enconman.2021.113877>.
- [23] Li J, Yang Z, Li H, Hu S, Duan Y, Yan J. Optimal schemes and benefits of recovering waste heat from data center for district heating by CO₂ transcritical heat pumps. *Energy Convers Manag* 2021;245:114591. <https://doi.org/10.1016/j.enconman.2021.114591>.
- [24] Zhou X, Xin Z, Tang W, Sheng K, Wu Z. Comparative study for waste heat recovery in immersion cooling data centers with district heating and organic Rankine cycle (ORC). *Appl Therm Eng* 2024;242:122479. <https://doi.org/10.1016/j.applthermaleng.2024.122479>.
- [25] Chen H, Peng Y, Wang Y. Thermodynamic analysis of hybrid cooling system integrated with waste heat reusing and peak load shifting for data center. *Energy Convers Manag* 2019;183:427–39. <https://doi.org/10.1016/j.enconman.2018.12.117>.
- [26] Haywood A, Sherbeck J, Phelan P, Varsamopoulos G, Gupta SK. Thermodynamic feasibility of harvesting data center waste heat to drive an absorption chiller. *Energy Convers Manag* 2012;58:26–34. <https://doi.org/10.1016/j.enconman.2011.12.017>.
- [27] Chen X, Pan M, Li X, Zhang K. Multi-mode operation and thermo-economic analyses of combined cooling and power systems for recovering waste heat from data centers. *Energy Convers Manag* 2022;266:115820. <https://doi.org/10.1016/j.enconman.2022.115820>.
- [28] Zhou X, Lin X, Su W, Ding R, Liang Y. Thermo-economic potential of Carnot batteries for the waste heat recovery of liquid-cooled data centers with different combinations of heat pumps and organic Rankine cycles. *Energies* 2025;18(6):1556. <https://doi.org/10.3390/en18061556>.
- [29] Poletto C, Pascale AD, Ottaviano S, Dumont O, Branchini L. Techno-economic assessment of a Carnot battery thermally integrated with a data center. *Appl Therm Eng* 2025;260:124952. <https://doi.org/10.1016/j.applthermaleng.2024.124952>.
- [30] Wang P, et al. District heating utilizing waste heat of a data center: high-temperature heat pumps☆. *Energ Build* 2024;315:114327. <https://doi.org/10.1016/j.enbuild.2024.114327>.
- [31] Corigliano O, Algieri A, Fragiaco P. Turning data center waste heat into energy: a guide to organic Rankine cycle system design and performance evaluation. *Appl Sci* 2024;14(14):6046. <https://doi.org/10.3390/app14146046>.
- [32] Zhu Y, et al. Optimal dispatch of a novel integrated energy system combined with multi-output organic Rankine cycle and hybrid energy storage. *Appl Energy* 2023; 343:121113. <https://doi.org/10.1016/j.apenergy.2023.121113>.
- [33] Chen S, et al. MILP optimization of the multi-heat pump waste heat recovery system integrated with full-free cooling data center through lake water. *Energy* 2025;318:134823. <https://doi.org/10.1016/j.energy.2025.134823>.
- [34] Takci MT, Qadrdan M, Summers J, Gustafsson J. Data centres as a source of flexibility for power systems. *Energy Rep* 2025;13:3661–71.
- [35] Lepour D, Loustau J, Terrier C, Maréchal F. REHO: a decision support tool for renewable energy communities. *J Open Source Softw* 2024;9(103):6734. <https://doi.org/10.21105/joss.06734>.
- [36] Transparency Platform. Accessed: May 25, 2025. [Online]. Available: <https://netransparency.entsoe.eu/market/energyPrices?appState=%7B%22sa%22%3A%5B%22BZN%7C10Y1001C-00098F%22%5D%2C%22st%22%3A%22BZN%22%2C%22mm%22%3Atrue%2C%22ma%22%3Afalse%2C%22sp%22%3A%22HALF%22%2C%22dt%22%3A%22MAP%22%2C%22df%22%3A%222025-05-25%22%2C%22tz%22%3A%22CET%22%7D>.
- [37] Reduce CO₂ Emissions with Real-Time Electricity Data. Emission. Accessed: May 25, 2025. [Online]. Available: <https://emission.io>.
- [38] Data center capital cost calculator. se.com. Accessed: May 26, 2025. [Online]. Available: <https://www.se.com/ww/en/work/solutions/system/s1/data-center-and-network-systems/trade-off-tools/data-center-capital-cost-calculator/>.
- [39] Prix d'une pompe à chaleur en suisse. Ias-engineering.ch. Accessed: May 26, 2025. [Online]. Available: <https://ias-engineering.ch/pac/prix-dune-pompe-a-chaleur-en-suisse/>.
- [40] Quel est le prix moyen des panneaux solaires en Suisse en 2024? helvetia-energy.ch. [Online]. Available: <https://www.helvetia-energy.ch/prix-moyen-panneaux-solaires-suisse/>.
- [41] Batteries solaires: les six questions les plus fréquentes. suisseenergie.ch. Accessed: May 26, 2025. [Online]. Available: <https://www.suisseenergie.ch/stories/batteries-solaires/>.
- [42] Devis pour un nouveau chauffage. scdi.ch. Accessed: May 26, 2025. [Online]. Available: <https://www.scdi.ch/fr/naturay-ultime-connecte-3.html>.
- [43] Chauffe-eau électrique sur socle Atlantic. scdi.ch. Accessed: May 26, 2025. [Online]. Available: <https://scdi.ch/fr/chauffe-eau-electrique-sur-socle-atlantic.html>.
- [44] 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 Sust Energ Rev* 2014;31:622–38. <https://doi.org/10.1016/j.rser.2013.12.007>.
- [45] Heat recovery and electricity generation – Heating Bits. Accessed: Jul. 09, 2025. [Online]. Available: <https://heatingbits.epfl.ch/research/heat-recovery-and-electricity-generation/>.
- [46] Zhou Yuekuan, Cao Sunliang. Quantification of energy flexibility of residential net-zero-energy buildings involved with dynamic operations of hybrid energy storages and diversified energy conversion strategies. *Elsevier* 2020;21:100304.
- [47] Jang S, Cho J. Waste heat recovery from data center cooling system for district heating. *ASHRAE Trans* 2024;130:1063–71.
- [48] Braz CG, Brigue R, Girardin L, Liu B, Maréchal F. Decomposition method to evaluate district heating/cooling network potential at urban scale. In: Book of abstract of the 34th european symposium on computer aided process engineering/15th international symposium on process systems engineering (ESCAPE34/PSE24), June 2–6, 2024, Florence, Italy. Associazione Italiana Di Ingegneria Chimica; 2024.
- [49] Cao Y, et al. Data-driven flexibility assessment for internet data center towards periodic batch workloads. *Appl Energy* 2022;324:119665.
- [50] Ghatikar G. Demand response opportunities and enabling technologies for data centers: Findings from field studies. 2012.
- [51] Pina EA, Lozano MA, Serra LM. A multiperiod multiobjective framework for the synthesis of trigeneration systems in tertiary sector buildings. *Int J Energy Res* 2020;44(2):1140–66. <https://doi.org/10.1002/er.5006>.
- [52] Fazlollahi S. Decomposition optimization strategy for the design and operation of district energy systems. Lausanne: STI; 2014. doi:<https://doi.org/10.5075/epfl-thesis-6130>.
- [53] Switzer J, Raghavan B. Information batteries: storing opportunity power with speculative execution. *SIGENERGY Energy Inform Rev* 2022;1(1):1–11. <https://doi.org/10.1145/3508467.3508468>.