Data-Center Waste Heat for District Energy: A Modeling Framework

Imran Aziz, Georgia Institute of Technology
Jung-Ho Lewe, PhD, Georgia Institute of Technology
Qixun Li, Georgia Institute of Technology

ABSTRACT

Large data centers reject steady, low-grade heat. Much of this heat is currently wasted, while nearby buildings need hot water and seasonal space heating. This project develops a clear, design-oriented modeling framework to reuse data-center heat in a district energy network serving a surrounding "data-center village." The study evaluates a simplified district-heating connection in which recovered data-center heat is upgraded and supplied to a representative building heating load. A plant model built in Modelica compares three system configurations: air cooling with heat recovery (Air + HRC), direct liquid cooling with heat recovery (DLC + HRC), and DLC with heat recovery and thermal energy storage (DLC + HRC + TES). The model represents rack-side heat capture, hydronic loops, plate heat exchangers, and heat-recovery chillers (HRCs) that raise temperatures to approximately 54–60 °C for district use, while providing simultaneous cooling back to the data hall. Basic hydraulics (temperature difference, flow, head) and representative receiver-building heating demand represented as an aggregate hourly load are included to show the system-level balance. Results highlight a simple insight: warmer rack return from liquid cooling improves the temperature quality and recoverability of the available heat and reduces the lift required from the HRC. Within the modeled cases, this supports greater heat delivery under the prescribed receiver-load conditions; TES heating benefits are evaluated during two representative winter periods. The framework helps engineers compare cooling configurations and evaluate key system components to deliver meaningful heat reuse while meeting the modeled data-center cooling load.

Introduction

The rapid expansion of cloud computing, artificial intelligence (AI), and digital infrastructure has led to a sustained increase in global data center capacity. As a result, data centers have become an increasingly important component of the global energy system. Recent analyses estimate that data centers account for approximately 1–2% of global electricity consumption, with projected growth driven largely by high-performance computing and AI workloads (International Energy Agency 2023; Masanet et al. 2020). Nearly all electrical energy consumed by information technology (IT) equipment is ultimately dissipated as heat. Consequently, large data centers represent substantial and relatively continuous sources of thermal energy that must be removed to ensure reliable operation of computing hardware.

Thermal management plays a critical role in maintaining acceptable operating temperatures for electronic components at the chip, server, rack, and room levels. Elevated temperatures can degrade semiconductor performance, reduce system reliability, and increase failure rates (ASHRAE 2021). Cooling infrastructure therefore constitutes a significant portion of total facility energy consumption in modern data centers. Traditional cooling architectures rely primarily on air-based systems in which computer room air conditioners (CRAC) or computer room air handlers (CRAH) circulate conditioned air through hot-aisle and cold-aisle configurations. While these systems remain widely deployed, increasing rack power densities have exposed the limitations of air cooling due to the relatively low heat capacity and thermal conductivity of air. To support higher computing densities, liquid-assisted cooling technologies are increasingly being adopted. These include rear-door heat exchangers (RDHx) and direct-to-chip liquid cooling (DLC), where coolant loops remove heat directly from high-power components such as CPUs and GPUs. By capturing heat closer to the source and utilizing fluids with higher thermal capacity than air, these systems enable more efficient thermal management and allow higher rack power densities (Li et al. 2024; Zimmermann et al. 2012). Other advanced approaches, such as immersion cooling, have also been explored in the literature for high-density applications (Zhang et al. 2022). Historically, most data center cooling systems rejected heat to the atmosphere through cooling towers or dry coolers. Because the temperature of the rejected heat was relatively low compared with conventional heating system requirements, this energy has often been considered waste heat. However, recent advances in industrial heat pumps and heat recovery chillers (HRCs) now enable low-grade heat to be upgraded to useful temperature levels for space heating and domestic hot water applications. These developments have stimulated growing interest in utilizing data center waste heat within district heating systems and integrated energy networks. Recent literature highlights the growing role of advanced liquid cooling technologies in enabling effective data center waste heat recovery for district heating applications. For instance, Cho and Moon (2025) demonstrate that direct-to-chip cold plate systems can achieve heat capture ratios of up to 94%, indicating strong potential for recovering high-grade heat suitable for reuse. Complementing this, Azarifar, Arik, and Chang (2024) show that liquid-to-liquid (L2L) and liquid-to-air (L2A) architectures provide robust thermal stability under varying loads. Additional studies show that coolant flow rate, inlet temperature, and fluid properties strongly influence thermal performance (Cheng et al. 2020; Taddeo et al. 2023). From an operational perspective, Gheni, Kerskes, and Stergiaropoulos (2026) highlight the importance of dynamic supply water temperature adjustment for improving efficiency under varying conditions. At the system level, Stein and Gill (2024) show that combining heat recovery chillers with thermal storage improves flexibility and supports building electrification. However, challenges related to integration complexity and scalability remain key barriers to widespread deployment (Azarifar, Arik, and Chang 2024).

Despite growing interest in data center heat reuse, several challenges remain. A key issue is the temperature level of recoverable heat, which depends on the cooling architecture. Higher return temperatures reduce the lift required from heat recovery chillers and improve efficiency. Another challenge is the mismatch between steady heat generation from data centers and variable building heating demand, which limits how much heat can be effectively utilized. Thermal energy storage (TES) can help address this mismatch by storing excess heat during low-demand periods and releasing it when demand increases (Rezaie and Rosen 2012; Guelpa and Verda 2019). However, its role in data center heat recovery systems, especially across different cooling configurations, remains less explored. Many existing studies also lack a system-level perspective, often focusing on either cooling performance or district heating independently rather than their interaction.

To address this gap, this study develops a Modelica-based framework to evaluate how a data center can serve as a heat source for nearby buildings. The model represents cooling systems, heat recovery chillers (HRCs), thermal storage, and building heating demand, capturing how recovered heat is used, stored, or rejected under different operating conditions. Three configurations are compared: air cooling with heat recovery (Air + HRC), direct liquid cooling with heat recovery (DLC + HRC), and DLC combined with both heat recovery and TES (DLC + HRC + TES). These are used to evaluate system performance, energy efficiency, and heat reuse potential. The main contributions of this work are as follows. First, a system-level framework is developed to evaluate air and liquid cooling with heat recovery within a district energy context. Second, TES is incorporated to address the mismatch between heat generation and demand, enabling targeted assessment during representative winter operating periods. Third, the study quantifies how heat recovery and TES affect system efficiency and the ability to meet building heating demand.

Methodology:

This study develops a system-level modeling framework to evaluate the potential for recovering waste heat from data center cooling systems and integrating it with a district energy system. Two cooling architectures are evaluated: conventional air cooling and direct-to-chip liquid cooling (DLC). Both architectures are coupled to heat recovery in the comparative cases, while TES is added only to the third configuration; the three modeled cases are Air + HRC, DLC + HRC, and DLC + HRC + TES. Each approach is modeled to represent heat capture and cooling performance, and the results are incorporated into a system-level Modelica model that includes heat recovery chillers (HRCs), thermal energy storage (TES), and district heating loads. The methodology combines operational data with system-level dynamic modeling to assess system performance under varying operating conditions.

Air-Cooled Data Center Cooling Model:

The air-cooling architecture represents a conventional data center in which server heat is removed by computer room air handling (CRAH) units connected to a centralized chilled-water plant. Conditioned air is supplied to the server racks and returned to the CRAHs, where heat is transferred through cooling coils to the chilled-water loop. The plant contains six CRAH units, represented at an aggregate level through the total cooling load. The air-cooling plant, together with its heat-recovery and district-heating interfaces, is shown in Figure 1.

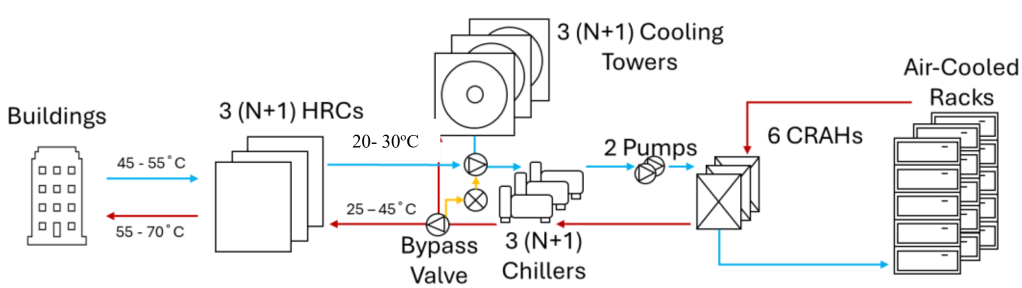

Figure 1. Schematic of air-cooled data center, heat recovery, and district heating integration.

The annual variation of IT-equipment electrical load and the corresponding prescribed cooling load is shown in Figure 2. The IT-equipment electrical load, ranging from approximately 900 kW to 1350 kW, provides the denominator for the modeled PUE calculation, while the cooling-load series (800–1170 kW) is the thermal boundary condition applied to the cooling models. The two series are paired hourly but are not treated as identical quantities. This variation reflects changes in computing activity and operating conditions. Figure 3 presents the corresponding average daily Atlanta wet-bulb temperature and modeled heating demand over the simulation year. The model represents a representative data-center floor with an approximately 1.5 MW electrical capacity. Its hourly IT-equipment electrical load ranges from approximately 0.9 MW to 1.35 MW over the simulated year. This load serves as the basis for the cooling, heat-recovery, and thermal-storage analyses. These ranges were selected from the hourly operating profile of the modeled data-center floor within an approximately 10 MW campus facility. The selected floor has an electrical capacity of roughly 1.5 MW, and its IT-equipment electrical load varied between about 0.9 and 1.35 MW over the study year. The corresponding 0.8–1.17 MW cooling requirement is the prescribed thermal load associated with that same hourly profile; the floor was selected to provide a single, capacity-consistent system for comparison and the results may be scaled cautiously to larger installations.

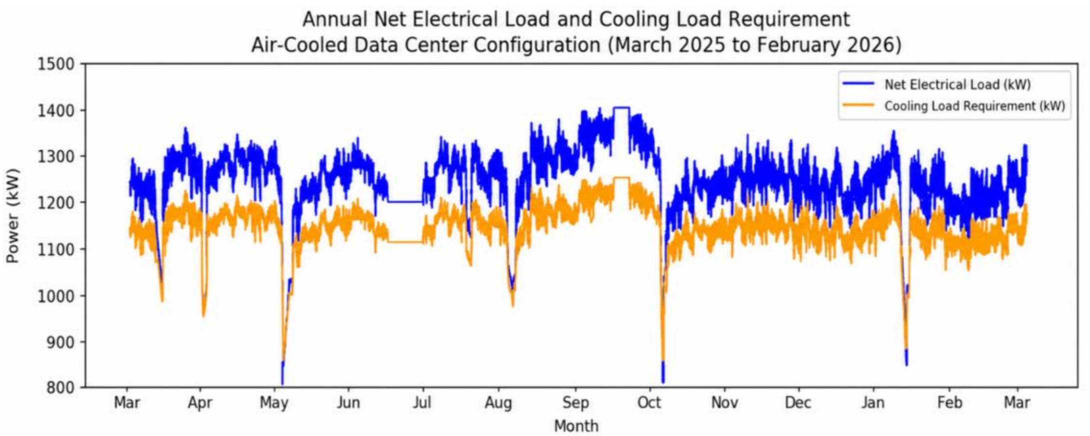

Figure 2. Annual variation of IT-equipment electrical load and corresponding cooling load requirement for the baseline air-cooled configuration from March 2025 to February 2026.

The six CRAHs are represented at the aggregate plant level. For the comparable 2025 period, the modeled average CRAH supply and return air temperatures were 17.06°C and 28.31°C, respectively, compared with measured averages of 18.68°C and 27.73°C across the six Level 1 CRAHs. The modeled return temperature is therefore within approximately 0.58°C of the measured value, indicating that the model reasonably represents the return-side thermal condition. Accordingly, the broader supply and return ranges of 13.6–19.5°C and 18.3–30.9°C are retained as the modeled annual operating envelope rather than presented as measured ranges. This comparison is an aggregate mean-temperature check rather than a full hourly dynamic validation.

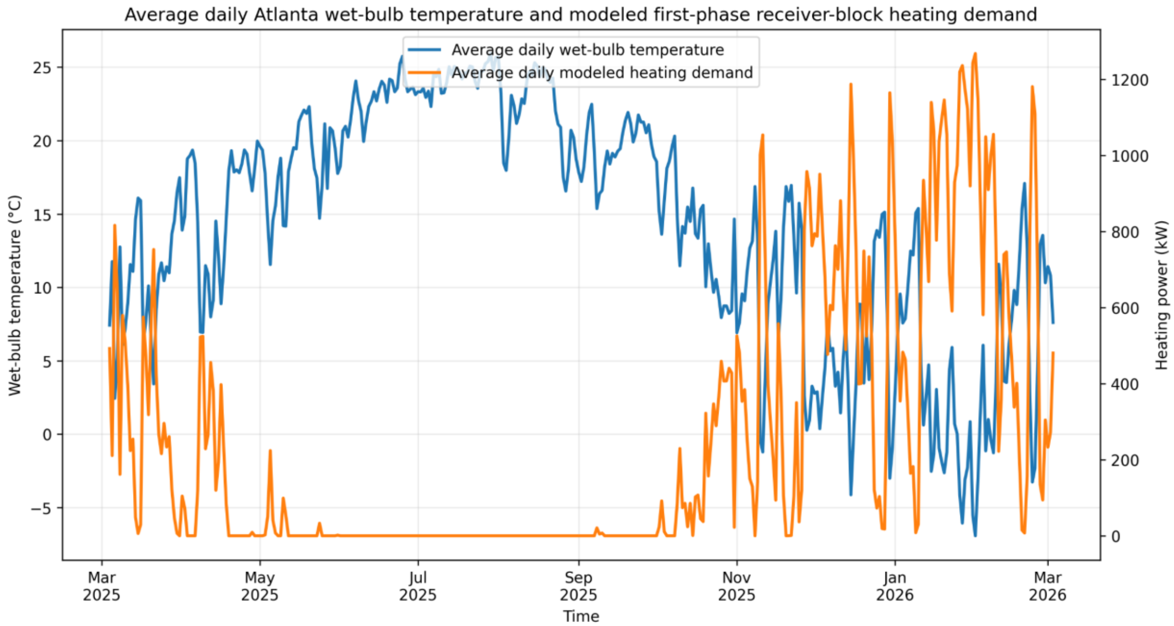

Figure 3. Average daily Atlanta wet-bulb temperature and modeled heating demand for the first-phase receiver block during the annual simulation period (March 2025 to February 2026). Wet-bulb temperature is shown because it drives cooling-tower performance; the receiver heating-demand profile is prescribed separately.

The chilled-water plant includes two water-cooled centrifugal chillers operating in parallel, with one additional unit providing N+1 redundancy. Each chiller is represented using a reformulated electric input ratio (EIR) model with a nominal capacity of approximately 760 kW. Nominal evaporator and condenser mass-flow rates remain approximately 21 kg/s and 23 kg/s per chiller, respectively. The chilled-water loop operates with supply temperatures of approximately 6–10°C and return temperatures of 12–17.5°C, corresponding to an evaporator temperature difference of about 5–8°C. Chilled-water flow per chiller varies from approximately 18 to 26 kg/s, with two parallel pumps operating and one pump on standby; the modeled pump head ranges from 30 to 55 ft (9–17 m). Condenser-water supply to the chillers ranges from approximately 24 to 32°C, while return water leaving the chiller condensers ranges from approximately 30 to 38°C under base heat-rejection operation. Condenser-water flow varies from approximately 20 to 31 kg/s per chiller. Cooling-tower performance is represented using wet-bulb approach and variable-speed fan control, and a plate heat exchanger provides water-side economizer operation when ambient conditions allow the chilled-water supply requirement to be met without compressor operation. These air-side and water-side operating conditions apply to the air-cooling topology and are not imposed on the DLC configuration.

For the modeled Air + HRC configuration, the existing CRAH chilled-water plant and N+1 chiller arrangement are retained, and the heat recovery chiller is added as a condenser-water-side retrofit. Warm condenser water leaving the base chillers is routed to the HRC evaporator, the cooling towers, or both. The modeled HRC source-water envelope is approximately 30–45°C entering the evaporator and 20–35°C leaving it. The cooled stream returns toward the base-chiller condenser inlet, while the cooling towers and bypass path reject heat that cannot be recovered and maintain the required entering-condenser-water temperature.

This arrangement preserves the existing CRAH airflow, chilled-water flow, and chiller-staging assumptions; the HRC changes the condenser-water heat-recovery path rather than the air-side mass-flow model. The base chillers and HRCs are represented using reformulated electric-input-ratio (EIR) Modelica models parameterized with manufacturer performance curves and available field data. Available capacity and compressor power vary with source-side temperature, load-side leaving-water temperature, and part-load ratio, and operation is restricted to the applicable performance envelope.

System performance is evaluated hourly over a full year. For all annual and selected week comparisons, modeled Power Usage Effectiveness (PUE) is defined as

$$PUE = \frac{P_{IT} + P_{cooling} + P_{other}}{P_{IT}}$$

where the denominator is the hourly IT-equipment electrical load. The numerator is the denominator plus the electrical consumption of the explicitly represented equipment: base-chiller and/or HRC compressors, modeled hydronic circulation pumps (including TCS/CDU and FWS pumping in the DLC cases), heat-rejection fans, TES-loop electrical parasitics where applicable, and a lumped auxiliary-load allowance representing power-distribution/PDU losses, humidity-control loads, and other facility electrical loads. The CRAHs are represented thermally using aggregate airflow and supply and return air-temperature boundary conditions; CRAH fan electricity is not explicitly simulated. The same calculation procedure, hourly IT-load denominator, and common auxiliary-load assumptions are applied to all three configurations, with case-specific equipment power included only where it is explicitly represented. Accordingly, the reported modeled PUE is intended for relative comparison within the modeled electrical boundary rather than as a complete facility-level PUE prediction. PUE is retained as a data-center electrical metric and does not credit useful heat exported to the receiver block. Broader system benefits should be evaluated separately using delivered heat, heat-recovery and pumping electricity, avoided conventional heating energy, distribution losses, and building-side equipment requirements. The present analysis evaluates the data-center cooling, heat-recovery, thermal-storage, and heat-delivery interface rather than a complete building-by-building HVAC replacement analysis. Building cooling demand and chilled-water export to the receiver buildings are not modeled in the present study.

Direct-to-Chip Liquid Cooling Model:

Direct-to-chip liquid cooling (DLC) represents a direct-to-chip liquid-cooling architecture in which liquid coolant is circulated through cold plates attached directly to high-power electronic components such as CPUs and GPUs. The coolant absorbs heat at the chip level and transfers it to a secondary facility water loop through a heat exchanger. In this paper, the rack-side coolant loop is referred to as the technology cooling system (TCS), while the secondary plant-side loop is referred to as the facility water system (FWS). At the campus facility, DLC was available for a limited number of racks using the CHx80 cooling system. For a capacity-consistent comparison with the air-cooling configuration, the modeled plant is scaled to 20 parallel CHx80 units, corresponding to approximately 1.6 MW of aggregate nominal CDU capacity; the hourly modeled liquid-cooling load remains below this installed nominal capacity. In the modeled DLC configurations, the complete prescribed cooling load is assigned to the TCS/CDU/FWS liquid pathway, and residual room-air cooling is not represented.

Operational measurements were used to define representative boundary conditions for the liquid-cooling loop. Specifically, measured supply and return coolant temperatures at the rack manifolds were used to characterize the temperature range over which heat is captured. Heat transfer within the liquid-cooling loop is calculated using the standard energy balance below, where $\dot{Q}$ is the transferred heat rate, $\dot{m}$ is the coolant mass-flow rate, c_p is the coolant specific heat, and T_{ret} and T_{sup} are the return and supply temperatures, respectively.

$$\dot{Q} = \dot{m}\, c_p \,(T_{ret} - T_{sup})$$

Compared to air cooling, liquid cooling benefits from the higher heat capacity and thermal conductivity of the working fluid, enabling heat to be captured at higher temperatures with lower thermal resistance (Zimmermann et al. 2012). Higher coolant return temperatures reduce the temperature lift required by heat recovery equipment, improving the efficiency of waste heat reuse. In the system model, the rack-level coolant loop is connected to a facility water loop through a plate heat exchanger. A key difference between the Air + HRC and DLC + HRC cases is the rack-side heat-capture path and its resulting temperature level. The plant topologies also differ: the Air + HRC case retains the base chiller and recovers heat from its condenser-water loop, whereas the DLC + HRC case supplies the HRC from the warmer FWS loop. The comparison therefore evaluates complete integration strategies rather than an isolated component-only air-versus-liquid substitution. The recovered heat is then directed either to heat rejection systems or to heat recovery chillers (HRCs), where it is upgraded for use in district heating applications. The full integration from rack-level heat capture to district heat recovery is illustrated in Figure 4. On the facility-water-system (FWS) side, the CDUs are connected to common supply and return headers. The modeled FWS supply temperature ranges from 20.0 to 30.0°C, with an annual average of 22.5°C, while the mass-flow-weighted FWS return ranges from 25.8 to 44.6°C, with an annual average of 36.1°C. Each CDU serves a local technology-cooling-system (TCS) rack loop; therefore, the reported TCS values are equivalent mass-flow-weighted temperatures rather than one physical common TCS header. The equivalent TCS supply ranges from 26.1 to 45.6°C, averaging 37.0°C, and the equivalent TCS return ranges from 33.3 to 64.3°C, averaging 54.1°C. The reported minima and maxima are pointwise extrema across the annual hourly series and should not be interpreted as persistent operating conditions. The 8.9 K (16°F) design temperature difference applies to the HRC evaporator/source loop downstream of the intermediate heat exchanger and is not imposed directly on either the FWS or TCS loop. The FWS and TCS temperatures remain functions of data-center cooling load, active CDU count, loop flow rates, and CDU approach temperature.

The available campus installation contains fewer CHx80 cooling distribution units (CDUs). For a capacity-consistent comparison, the modeled DLC plant was scaled to 20 parallel CHx80 units, providing approximately 1.6 MW of nominal CDU capacity. The hourly liquid-cooling load is divided among the operating CDUs, and the loop temperatures are calculated from the per-CDU heat load, primary and secondary flow rates, and the manufacturer load-dependent approach-temperature relationship. Scaling the number of parallel CDUs increases aggregate capacity and flow; it does not multiply the supply or return temperatures because heat load and flow scale proportionally according to the following relation.

$$\Delta T = \frac{\dot{Q}}{\dot{m}\, c_p}$$

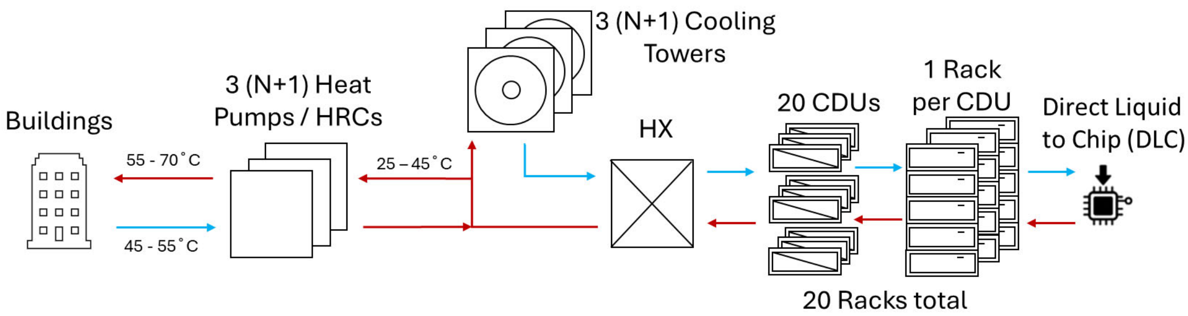

Figure 4. DLC schematic showing the heat-flow path from the racks through the liquid-cooling and heat-recovery systems to the receiver buildings.

System-Level Modelica Energy Model

A system-level model was developed using the Modelica modeling language to evaluate the interaction between data center cooling systems and district energy infrastructure. The model integrates the cooling architectures described above within a common framework that includes heat recovery chillers (HRCs), thermal energy storage (TES), and district heating loads.

The modeled data center represents a 0.8–1.17 MW cooling load, scaled from operational measurements obtained from a 10 MW campus data center. Air-cooling data was available at the system level, while DLC data was obtained from a limited number of racks equipped with CHx80 units. These measurements were scaled proportionally to represent approximately 20 CHx80 units, corresponding to a comparable system-level cooling capacity. The cooling loads are connected to facility water loops and heat rejection systems, including dry coolers and cooling towers. Several heat rejection and recovery configurations were evaluated, including direct heat rejection using dry coolers, cooling tower-based heat rejection, and the integration of heat recovery chillers producing hot water in the range of 54–60°C. To address the mismatch between continuous data center heat generation and variable building heating demand, the model incorporates thermal energy storage (TES) in the form of a stratified hot water tank connected to a simplified district heating network. The system-level Modelica implementation, including TES integration, is shown in Figure 5. The TES model represents a 160 m^3 stratified hot-water tank divided into four control volumes. Layer 1 is the top layer nearest the outlet, and Layer 4 is the bottom layer nearest the inlet.

The recovered heat is evaluated against a representative first-phase receiver block in the immediate vicinity of the data center. The block comprises four existing buildings with a combined floor area of approximately 203,500 GSF and a diversity-adjusted peak heating demand of approximately 1.31 MW. A representative hourly heating-demand profile for this block is modeled in Modelica and used as the receiving-demand input for the heat-recovery and thermal-storage simulations. The buildings are currently served through existing utility infrastructure and are not physically connected to the data-center plant; practical implementation would require underground low-temperature-hot-water piping, building-side heat exchangers, metering, and control integration. These distribution and capital-infrastructure requirements are outside the scope of the present study. The integrated model captures hydraulic flows, temperature evolution, and heat transfer interactions between data center cooling systems, heat recovery equipment, thermal storage, and district loads. Simulations were conducted under varying weather conditions and demand profiles to evaluate the fraction of recoverable heat and overall system performance under different operating scenarios. The prescribed receiver profile is used as a system-level demand input and is not presented as a calibrated building-by-building load model. The reported results are specific to the Atlanta climate and the representative receiver block. Colder climates and heating-dominant building mixes are expected to increase the opportunity for recovered-heat utilization, whereas warm climates or cooling-dominant districts may reduce annual heating utilization and require alternative heat sinks.

Accordingly, the two HRC arrangements represent different integration strategies. The air-based case is a condenser-water-side retrofit that retains the existing chillers, whereas the DLC case recovers heat directly from the warmer CDU facility-water loop through an intermediate heat exchanger. In the DLC case, the HRC extracts heat from the hydraulically separated source loop when useful heating or TES charging is available; otherwise, the cooling towers reject the remaining data-center heat. A direct chilled-water-side HRC for the CRAH plant would also be feasible but was not evaluated in this comparison.

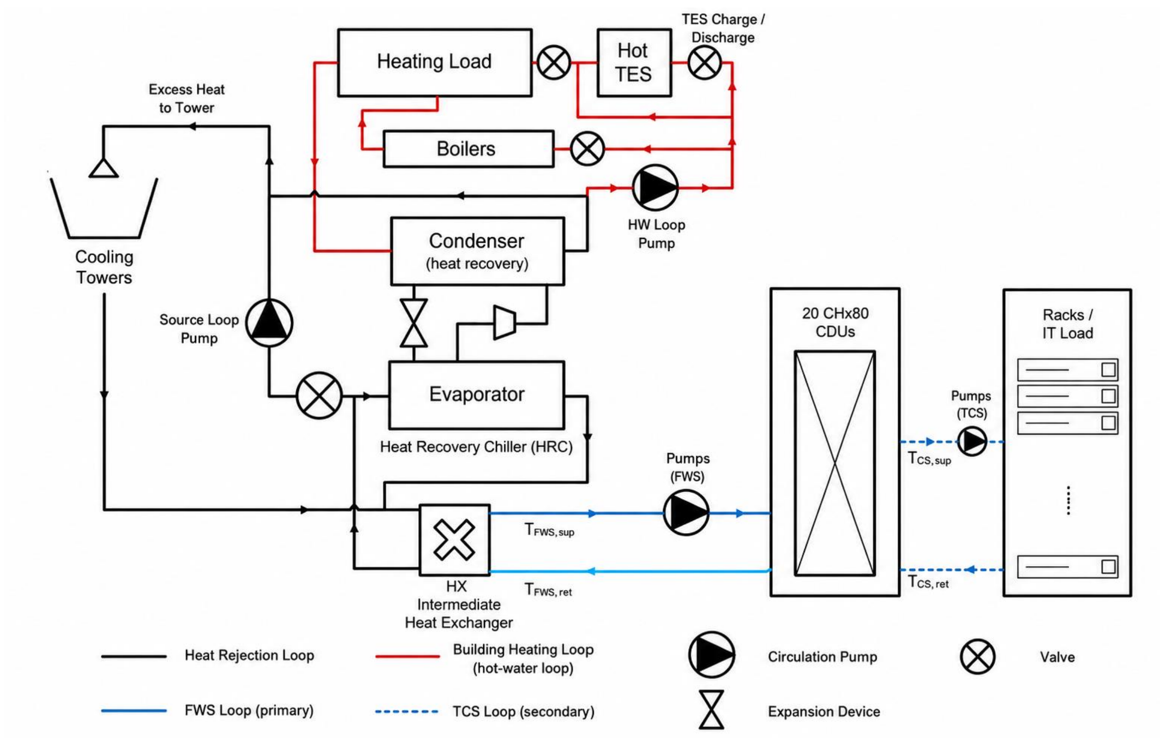

Figure 5. Modelica-based system schematic of the facility water system integrating technology-side (TCS) and facility-side loops with waste heat recovery using a heat-recovery chiller (HRC). The model also incorporates thermal energy storage (TES) and a dynamic building heating load to evaluate integrated cooling and heat recovery performance.

Results:

Building on the Modelica-based plant model described in the methodology section, the system performance is evaluated by comparing three cooling configurations: air cooling with heat recovery (Air + HRC), direct liquid cooling with heat recovery (DLC + HRC), and direct liquid cooling combined with both heat recovery and thermal energy storage (DLC + HRC + TES). The simulations capture interactions between the data center and the district energy network, allowing evaluation of both energy efficiency and heat recovery under different operating conditions. Over the annual simulation, both liquid cooling configurations consistently show lower and more stable modeled PUE than Air + HRC. Among the two, the DLC + HRC system achieves the lowest PUE for most of the year. Adding TES introduces small variations due to charging and discharging, but it does not significantly impact overall performance, as illustrated in Figure 6.

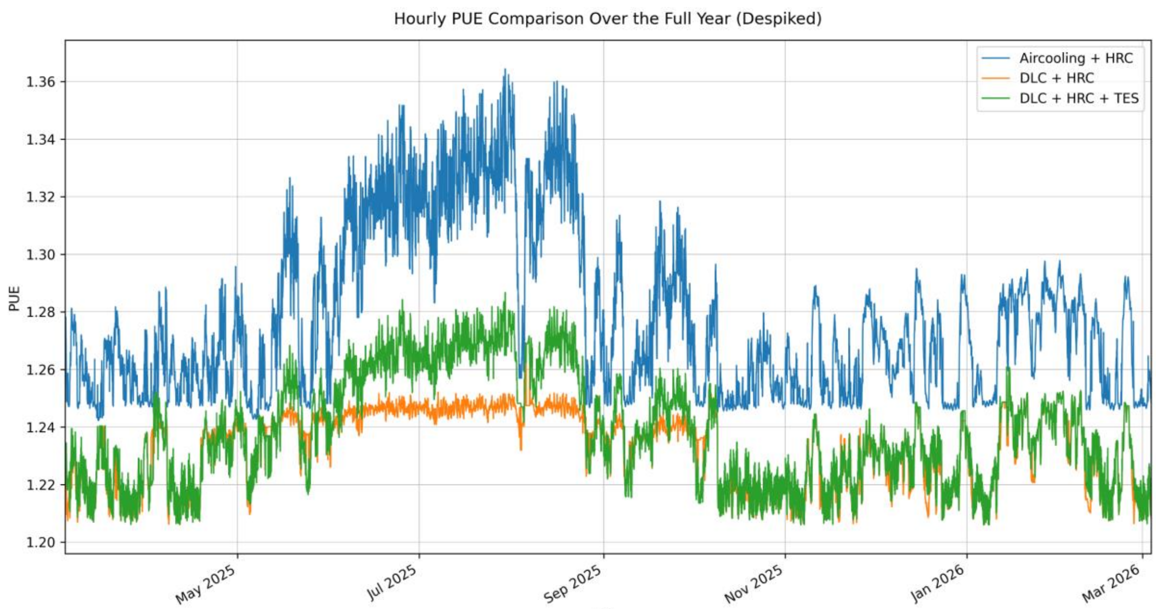

Figure 6. Annual comparison of modeled Power Usage Effectiveness (PUE) for Air + HRC and two direct liquid cooling configurations (DLC + HRC and DLC + HRC + TES). The Air + HRC configuration shows higher and more variable modeled PUE across the year, while both DLC configurations maintain consistently lower and more stable modeled PUE.

The annual comparison uses the same hourly IT-load and weather sequence for Air + HRC, DLC + HRC, and DLC + HRC + TES; PUE is reported as a dimensionless electrical-efficiency metric and does not credit exported thermal energy. The three configurations are compared using the same modeled electrical boundary defined in the methodology.

A reduction in modeled PUE from 1.34 to 1.26 corresponds to a decrease of 0.08 kW in modeled electrical demand per kilowatt of IT load. For a data center operating at 1 MW of IT load, this represents an approximately 80 kW reduction in modeled electrical demand, or about a 6.0% reduction relative to the 1.34-PUE case. Although PUE is useful for comparing data-center electrical efficiency, it does not account for the useful thermal energy delivered to the receiving buildings. Selected-period heat delivery and unmet building-heating demand are therefore evaluated separately. Seasonal trends help explain how TES behaves under different conditions. In summer, heating demand is very low, so TES remains mostly inactive, and the DLC + HRC system achieves the lowest modeled PUE. In winter, as heating demand increases, TES becomes active and supports the system, while both DLC configurations continue to show lower modeled PUE than Air + HRC. Although the TES case shows slightly higher PUE due to modeled TES-loop circulation/control parasitics and thermal losses, the overall impact on efficiency is small, as shown in Figure 7. The same PUE definition is used for the annual and selected-week comparisons. The annual PUE comparison is based on the full-year hourly simulation. The TES heating-benefit results presented below are targeted evaluations of two representative winter periods and should not be interpreted as an annual TES-coverage assessment.

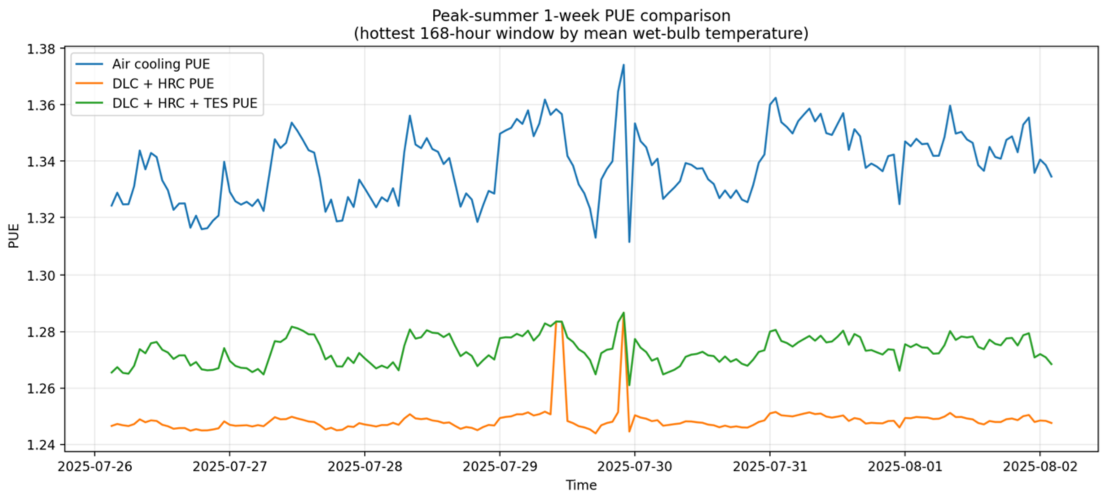

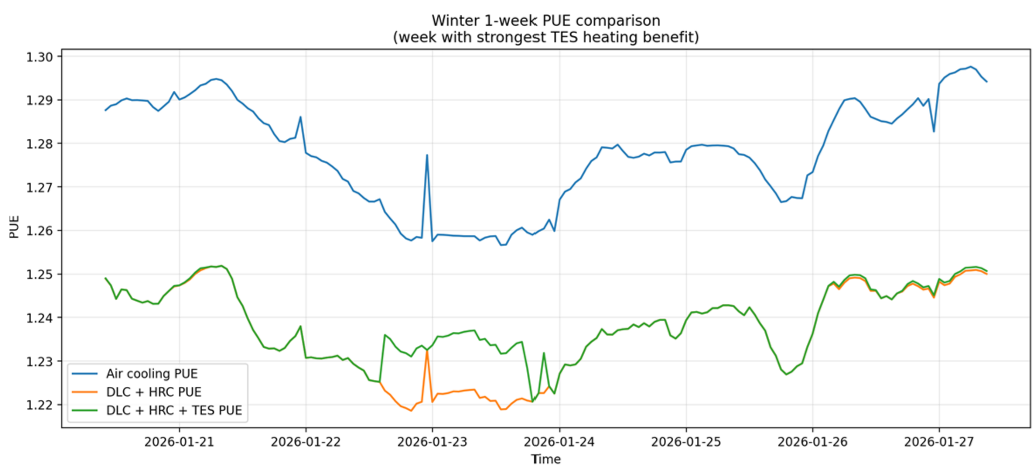

Figure 7. Weekly PUE comparison for peak summer and peak winter conditions for Air + HRC, DLC + HRC, and DLC + HRC + TES systems, illustrating seasonal variation in TES utilization and its impact on system performance.

Thermal Energy Storage (TES) Performance Under Winter Demand Conditions

The underlying thermal behavior of TES is characterized by stable cyclic operation with stratification retained within the four-control-volume tank model. Repeated charging and discharging cycles show limited modeled mixing between adjacent layers, indicating the system's ability to store and release heat effectively over time. This behavior supports the role of TES in enabling temporal balancing between heat recovery and demand, as observed in Figure 8.

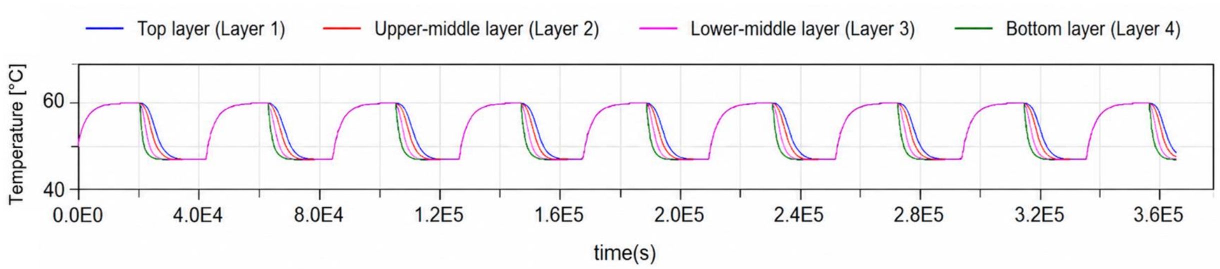

Figure 8. Temperature evolution in the stratified storage tank during repeated 6-hour charging and discharging cycles. The plot shows a 100-hour simulation in March, illustrating the thermal response of the tank layers and the preservation of stratification over successive cycles. The modeled TES has a total volume of 160 m³ and four control volumes representing the top, upper-middle, lower-middle, and bottom layers.

To demonstrate the effectiveness of thermal energy storage (TES), a representative winter week (January 20–27, 2026) was selected, during which the campus building experienced significant variability in heating demand, with an average demand of 950.1 kW and peak values reaching 1300 kW. This temporal fluctuation enabled periods where the heat recovered by the DLC + HRC system exceeded the instantaneous heating requirement, allowing excess thermal energy to be stored in TES. As shown in Figure 9, the DLC + HRC system alone is unable to consistently track the demand curve, resulting in noticeable shortfalls during peak hours. In contrast, the DLC + HRC + TES configuration closely follows the demand profile by discharging stored heat during these peak periods. Consequently, the total unmet heating energy is reduced from 4514.4 kWh to 520.7 kWh, corresponding to an 88.5% reduction in unmet heat. For this representative variable-demand week, the result shows that TES can be highly effective when there is a temporal mismatch between heat availability and demand, enabling recovery and reuse of thermal energy that would otherwise be rejected.

A second winter period (January 26–February 3, 2026) was analyzed to evaluate TES performance under sustained high heating demand conditions. During this week, the building exhibited a consistently elevated heating demand, with an average of 1165.6 kW and peaks up to 1300 kW. Unlike the previous case, the demand remained close to or above the maximum recoverable heat from the HRC system for most hours, limiting the availability of excess heat for storage. As illustrated in Figure 10, the DLC + HRC system frequently falls below the heating demand curve, indicating persistent unmet heating. The addition of TES improves performance by partially bridging this gap, with the DLC + HRC + TES curve tracking closer to the demand. However, due to the lack of sufficient surplus heat for charging, TES is unable to fully eliminate the deficit. The total unmet heating energy decreases from 14,457 kWh to 10,842 kWh, corresponding to a 25% reduction, but a significant shortfall remains. This indicates that under sustained high-demand conditions, system performance is constrained by heat recovery capacity rather than temporal mismatch, thereby limiting the effectiveness of TES.

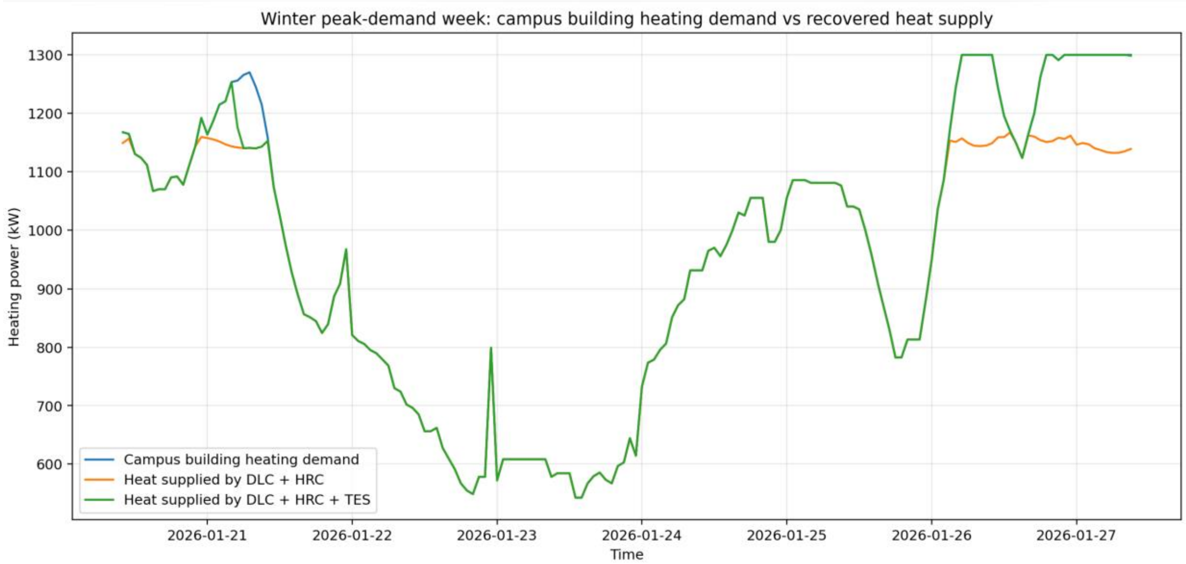

Figure 9. Winter week (January 20–27, 2026) comparison of campus heating demand, heat supplied by DLC + HRC, and heat supplied by DLC + HRC + TES. The figure shows that TES enables close tracking of peak heating demand by discharging stored thermal energy, significantly reducing unmet heating during high-demand periods.

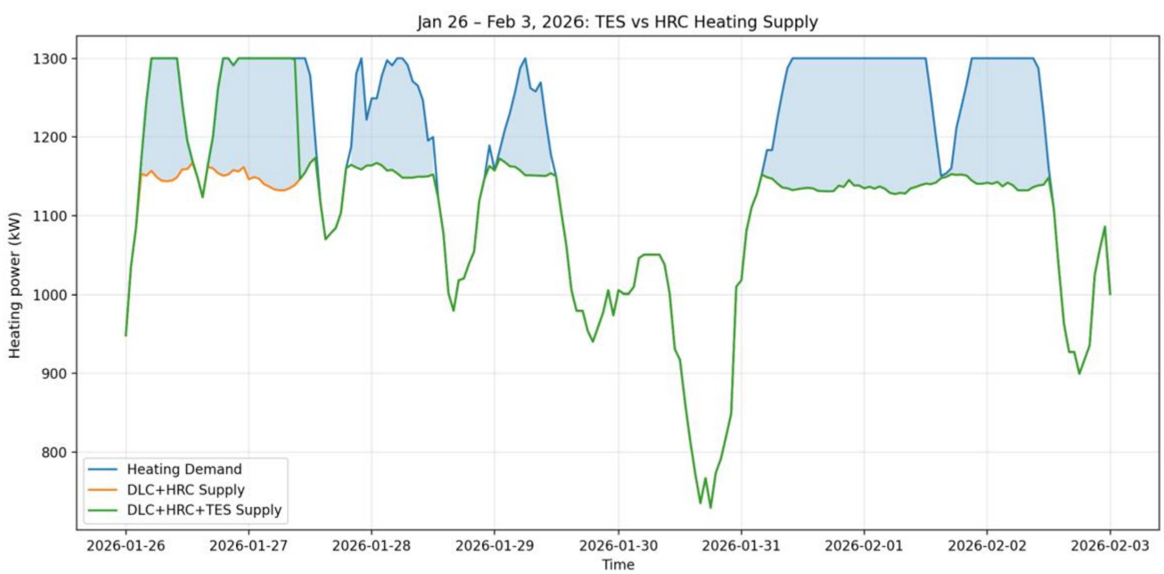

Figure 10. Winter period (January 26–February 3, 2026) comparison of heating demand and supplied heat for DLC + HRC and DLC + HRC + TES. Under sustained high-demand conditions, TES reduces but does not eliminate unmet heating due to limited availability of surplus recoverable heat for storage.

Overall, the results reinforce the central premise outlined in the abstract: data centers can act as potential heat sources within a district energy system when appropriate cooling and heat recovery configurations are employed. Figures 6 and 7 show that the DLC configurations provide lower modeled PUE than Air + HRC, while Figures 9 and 10 demonstrate how TES affects recovered-heat delivery during the selected winter periods. The addition of TES improves temporal matching between heat availability and building demand in the two selected winter periods, with the strongest reduction in unmet heating during the variable-demand week. However, the results also demonstrate that TES effectiveness depends on operating conditions, providing substantial benefits under fluctuating loads but limited gains under sustained high-demand scenarios where system performance becomes capacity-constrained. Within the modeled cases, higher rack return temperatures reduce HRC lift and improve the recoverability of the available heat under the prescribed receiver-load conditions while serving the same prescribed IT-load profile. Consequently, the integrated DLC + HRC and DLC + HRC + TES configurations provide a practical basis for evaluating data-center heat recovery and storage under the modeled district-heating load conditions.

Conclusions

This study presents a system-level modeling framework to evaluate the integration of data center cooling systems with district energy networks. By comparing air cooling with heat recovery (Air + HRC), direct liquid cooling with heat recovery (DLC + HRC), and DLC combined with thermal energy storage (DLC + HRC + TES), the results demonstrate the potential of data centers to serve as heat sources under the modeled operating and demand conditions for district heating applications.

Liquid cooling configurations consistently achieve lower and more stable modeled PUE in both the annual and selected-week comparisons than Air + HRC, while enabling higher-temperature heat recovery. The use of heat recovery chillers upgrades low-grade waste heat to usable levels (approximately 54–60°C), supporting the feasibility of heat export to district systems.

Thermal energy storage (TES) improves system performance by mitigating the temporal mismatch between heat generation and building demand. During the representative variable-demand winter week, TES significantly reduces unmet heating demand by up to 88.5% relative to DLC + HRC by shifting excess heat to peak-demand periods. However, under sustained high-demand conditions, its effectiveness is limited by the available heat recovery capacity rather than storage capability.

Overall, TES primarily enhances heat utilization and system flexibility rather than directly reducing electrical energy consumption. The proposed framework provides a practical basis for evaluating cooling configurations and heat recovery strategies in integrated data center–district energy systems.

Future Work

Future work should extend this framework to include immersion cooling, which has the potential to deliver higher coolant return temperatures and improved heat capture. This could further reduce heat recovery chiller lift and enhance overall system efficiency. Incorporating immersion cooling would enable a more complete comparison across cooling architectures and strengthen the assessment of heat reuse potential. Additional research should also focus on advanced control strategies for coordinating heat recovery, storage, and distribution under dynamic conditions, as well as incorporating economic analysis to support real-world deployment.

Future work should also investigate the impact of dynamic AI workload profiles on the district-heating-loop temperature setpoint and develop recommendations and risk-mitigation strategies for district-heating operators. These should include setpoint-reset limits, ramp-rate constraints, thermal-buffer and backup-heat requirements, forecast-based dispatch, and fail-safe controls for rapid workload changes or communication loss.

The results presented in the manuscript are specific to the Atlanta climate, the representative receiver block, and the modeled integration topologies. Future work will also evaluate multiple climate regions, alternative receiver-building compositions, different heating-water temperature requirements, building-load diversity, district-pipe losses, and the influence of seasonal load coincidence on heat-recovery and thermal-storage sizing. Future studies should also incorporate building cooling profiles and integrated district heating-and-cooling dispatch.

References

ASHRAE. 2021. Thermal Guidelines for Data Processing Environments. 5th ed. Atlanta, GA: American Society of Heating, Refrigerating and Air-Conditioning Engineers.

Azarifar, M., M. Arik, and J. Chang. 2024. “Liquid Cooling of Data Centers: A Necessity Facing Challenges.” Applied Thermal Engineering 247: 123112. https://doi.org/10.1016/j.applthermaleng.2024.123112.

Cheng, C.-C., P.-C. Chang, H.-C. Li, and F.-I. Hsu. 2020. “Design of a Single-Phase Immersion Cooling System through Experimental and Numerical Analysis.” International Journal of Heat and Mass Transfer 160: 120203. https://doi.org/10.1016/j.ijheatmasstransfer.2020.120203.

Cho, J., and J. Moon. 2025. “Numerical Coupling of Energy Efficiency and Thermal Performance for Cold Plate Cooling Optimization in High-Density Compute AI Data Centers.” Energy and Buildings: 116441. https://doi.org/10.1016/j.enbuild.2025.116441.

Gheni, M., H. Kerskes, and K. Stergiaropoulos. 2026. “Operational Analysis of the Cooling System in a Direct Liquid-Cooled Data Center: A Measurement and Simulation Study on the Impact of Supply Water Temperature.” Applied Energy 403: 127061.

Guelpa, E., and V. Verda. 2019. “Thermal Energy Storage in District Heating and Cooling Systems: A Review.” Applied Energy 252: 113474.

International Energy Agency (IEA). 2023. “Data Centres and Data Transmission Networks.” Paris: International Energy Agency.

Li, X., M. Li, Y. Zhang, Z. Han, and S. Wang. 2024. “Rack-Level Cooling Technologies for Data Centers: A Comprehensive Review.” Journal of Building Engineering 90: 109535. https://doi.org/10.1016/j.jobe.2024.109535.

Masanet, E., A. Shehabi, N. Lei, S. Smith, and J. Koomey. 2020. “Recalibrating Global Data Center Energy-Use Estimates.” Science 367 (6481): 984–986. https://doi.org/10.1126/science.aba3758.

Rezaie, B., and M. A. Rosen. 2012. “District Heating and Cooling: Review of Technology and Potential Enhancements.” Applied Energy 93: 2–10.

Stein, J., and B. Gill. 2024. “The Solution to Large Building Electrification: Heat Recovery Chillers and Condenser Water Storage.” In ACEEE Summer Study on Energy Efficiency in Buildings.

Taddeo, P., J. Romaní, J. Summers, J. Gustafsson, I. Martorell, and J. Salom. 2023. “Experimental and Numerical Analysis of the Thermal Behaviour of a Single-Phase Immersion-Cooled Data Centre.” Applied Thermal Engineering 234: 121260. https://doi.org/10.1016/j.applthermaleng.2023.121260.

Zhang, Y., Y. Zhao, S. Dai, B. Nie, H. Ma, J. Li, Q. Miao, Y. Jin, L. Tan, and Y. Ding. 2022. “Cooling Technologies for Data Centres and Telecommunication Base Stations: A Comprehensive Review.” Journal of Cleaner Production 334: 130280. https://doi.org/10.1016/j.jclepro.2021.130280.

Zimmermann, S., I. Meijer, M. K. Tiwari, S. Paredes, B. Michel, and D. Poulikakos. 2012. “Aquasar: A Hot Water Cooled Data Center with Direct Energy Reuse.” Energy 43 (1): 237–245. https://doi.org/10.1016/j.energy.2012.04.037.

AI Statement

AI-assisted tools (ChatGPT) were used for language refinement, rephrasing, and generating plots based on simulation data produced using the Modelica model. The results presented in this paper were obtained from Modelica simulations and were further processed using Python for data extraction and visualization. AI tools were not used to generate simulation results or develop the underlying numerical methodology or Modelica implementation. Technical interpretations, manuscript revisions, and final conclusions were reviewed and approved by the authors.