EDGE: Microgrid Data Center with Mixed Energy Storage

Rickard Brännvall*[†]
RISE ICE Datacenter
Luleå, Sweden
rickard.brannvall@ri.se

Mikko Siltala*
RISE ICE Datacenter
Luleå, Sweden
mikko.siltala@ri.se

Jonas Gustafsson*
RISE ICE Datacenter
Luleå, Sweden
jonas.gustafsson@ri.se

Jeffrey Sarkinen*
RISE ICE Datacenter
Luleå, Sweden
jeffrey.sarkinen@ri.se

Mattias Vesterlund*
RISE ICE Datacenter
Luleå, Sweden
mattias.vesterlund@ri.se

Jon Summers*
RISE ICE Datacenter
Luleå, Sweden
jon.summers@ri.se

ABSTRACT

Low latency requirements are expected to increase with 5G telecommunications driving data and compute to EDGE data centers located in cities near to end users.

This article presents a testbed for such data centers that has been built at RISE ICE Datacenter in northern Sweden in order to perform full stack experiments on load balancing, cooling, microgrid interactions and the use of renewable energy sources. This system is described with details on both hardware components and software implementations used for data collection and control. A use case for off-grid operation is presented to demonstrate how the test lab can be used for experiments on edge data center design, control and autonomous operation.

CCS CONCEPTS

• **Computer systems organization** → *Embedded systems*; • **Hardware** → **Enterprise level and data centers power issues**.

KEYWORDS

Data centers, Edge, Monitoring, Microgrid, Batteries, Thermal Energy Storage

ACM Reference Format:

Rickard Brännvall, Mikko Siltala, Jonas Gustafsson, Jeffrey Sarkinen, Mattias Vesterlund, and Jon Summers. 2020. EDGE: Microgrid Data Center with Mixed Energy Storage. In *The Eleventh ACM International Conference on Future Energy Systems (e-Energy'20)*, June 22–26, 2020, Virtual Event, Australia. ACM, New York, NY, USA, 8 pages. <https://doi.org/10.1145/3396851.3402656>

1 INTRODUCTION

The amount of data generated, processed, stored and transmitted has grown tremendously during the last decades, due to an increased digitalization of society and new digital services, including

streaming media, social media, IoT etc. To cope with the increased amount of data, the number of very large data centers has grown. This has helped to move computation and storage from local data centers to central hyperscale, often very efficient data centers providing the majority of cloud services used today. These large data centers are often located where there are good conditions for operating data centers, for example stable and environmental power supply, affordable land and good network infrastructure. However, these conditions are often not present in major cities, where people tend to live and work.

The latency introduced using centralized data centers (~20-30ms) has been considered acceptable as earlier mobile communication technologies (1/2/3/4G) operate with higher latencies making 20-30ms imperceptible. With 5G, latency decreases to 1-10 ms, meaning that the larger part of system latency is now occurring between the base station and the data center.

To enable full functionality for low-latency devices and latency-sensitive applications (eg., control loops, real time analytics), compute resources traditionally found in data centers need to move closer to the access points and end users, to the edge of the network. These future data centers, often referred to as edge data centers will probably be co-located with access points or other mobile network equipment. A grand challenge with the edge-concept is to provide these edge data centers with power and efficient heat rejection (cooling) systems. This is even more challenging than with traditional "rural" data centers, since the edge data centers will mostly be located in cities where the power availability is sparse, and heat rejection might be even trickier as the data center can be co-located inside existing buildings.

To cope with sparse energy supply, local electricity production and energy storage, both electrical and thermal, can be a way to deploy these data centers even if the grid power is not designed to support the dimensioned load of the data center(s). A microgrid solution could support both the data center operation and the power grid, through load balancing functionality.

In order to perform full stack experiments involving a 5G accessible, microgrid enabled data center, an EDGE test lab was set up at the RISE ICE Datacenter test bed in Luleå, Sweden. The lab provides opportunities for holistic experiments involving both the facility domain, including power and cooling, and the IT, as well as how these should co-operate to achieve the best possible system performance. It also allows for power load balancing and off-grid

*Research Institutes of Sweden, Digital Systems, Computer Science, ICE Datacenter

[†]Luleå University of Technology, Department of Computer Science, EISLAB

This work is licensed under a Creative Commons Attribution International 4.0 License.

e-Energy'20, June 22–26, 2020, Virtual Event, Australia

© 2020 Copyright held by the owner/author(s).

ACM ISBN 978-1-4503-8009-6/20/06.

<https://doi.org/10.1145/3396851.3402656>

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

operations, since energy storage and local electricity production using photovoltaics are available.

The RISE EDGE lab also includes thermal energy storage in the form of a cold tank, which set it apart from previous reported similar testbeds, see e.g. pioneers like Parasol [6] that focus on IT-load scheduling, use batteries and renewable energy sources - a previous study that emphasizes the important experience that can be gained from actually building a complete experimental data center over relying on simulations only.

This paper presents a complete system description and a first use case demonstration of how the EDGE data center lab can be used to evaluate different control strategies.

2 EXPERIMENTAL SET-UP

The EDGE lab is an experiment testbed at the RISE ICE Datacenter research facility set up with the purpose to explore design, control and operational considerations for small data centers, particularly around matters relating to energy efficiency, cost, self-sufficiency and autonomy.

2.1 System description

A photo of the system is presented in Figure 1, where most of the system components are visible. A schematic drawing of the electrical and thermal components is also presented in Figure 2. These components are further explained in the paragraphs below.

2.1.1 Data center module. The container module, (1) in Figure 1, is a fully contained data center, which has cold (front) and hot (back) aisle separation. Between the aisles are two server racks, and a computer room air handler (CRAH) unit above the racks. The module is constructed in sheet metal and is suitable for both indoor and outdoor use.

2.1.2 IT equipment. The IT equipment is installed in the two 42U server racks. Each rack contains 38 Dell PowerEdge R430 servers,

as well as 2 Dell N1548 switches, located at the top of the rack. The total idle load of the servers is about 6 kW, and the maximum load is roughly 12 kW.

Power is supplied to the equipment using Schleifenbauer rack PDUs which have network connectivity and metering capability at the inputs and outlets. The servers each contain one PSU with a nameplate power rating of 550W.

2.1.3 Computer room air handler. The overhead CRAH unit contains a fan which forces air from the hot aisle through an air-to-water heat exchanger to the cold aisle. The fan speed is controlled to keep a constant differential pressure between the two aisles. Adjustable dampers are used to direct the discharge air downwards towards the front of the servers.

2.1.4 Thermal Energy Storage tank. The thermal energy storage tank has a capacity of 2000 liters and contains water that is used to cool the module. Water flows out from the base of the tank, passes through the CRAH and returns to the top of the tank. This creates thermal stratification and a temperature gradient inside the tank where cold water is located at the bottom, and hot at the top. The hot water from the top of the tank is also pumped out to the refrigeration unit, and is returned cooler to the base of the tank. The tank and piping are insulated to reduce heating from ambient air.

Flow diffusers and nozzles. The inputs to the storage tank have flow diffusers that reduce the flow velocity inside the tank to avoid mixing the water and disturb the temperature gradient. Likewise the outputs have nozzles that increase the flow velocity in the pipes. These can be seen in the system picture as the red cone-shaped fittings attached to the tank.

Secondary user side pump. The secondary user side pump is used to create the water flow through the CRAH. The pump speed can be controlled, and is used to regulate the cold aisle temperature.

Figure 2: A schematic drawing of the facility components in the EDGE data center setup.

Figure 3: A schematic drawing of the chiller internal components in the EDGE data center setup.

Secondary shunt valve. The CRAH supply water temperature can be controlled using the secondary shunt valve by mixing cold water from the tank and warmer water from the CRAH return line.

2.1.5 Chiller. A Bluebox Tetris W Rev FC/NG chiller unit sits between the primary and secondary cooling loops, which use water and a glycol-water mixture respectively as fluid media. (To be precise the second loop uses a 30% ethylene glycol and water mixture, which for brevity will be denoted just as glycol going forward.)

The chiller unit consists of a vapor-compression refrigeration cycle and a free cooling heat exchanger and directs the flow of glycol to use the appropriate cooling method using a 3 way valve, shown in Figure 3. In addition to the (pure) chiller mode and the (pure) free cooling mode, it can also operate in partial free cooling mode where both methods are utilized.

Chiller mode. The chiller mode uses the vapor-compression cycle, the components of which are illustrated in Figure 3. The coolant liquid is vaporized to absorb heat from the warm water in the evaporator heat exchanger, and then compressed to remove heat to the glycol in the condenser heat exchanger.

Free cooling mode. When the outside temperature is low, the vapor-compressor cycle is not required to increase the temperature difference between the water and glycol, and the free cooling heat exchanger is used to directly transfer heat from water to glycol.

Partial free cooling mode. In this mode the 3-way valve is modulated to split the glycol flow to both the free cooling heat exchanger and the (now in reduced mode) vapor-compression cycle.

Pumps. The chiller unit contains two pumps to regulate the flow of the coolant water and glycol through the heat exchangers.

2.1.6 Cooling tower pipes. Insulated pipes carry the glycol from the chiller to the drycooler and back again.

Primary shunt valve. A shunt valve is used to regulate the temperature of glycol entering the chiller unit. When the outside temperature is below freezing the glycol returned from the drycooler is proportionally mixed with that which exits in the chiller so that

it can not freeze the water in the free cooling heat exchanger, or any water that condenses inside the chiller unit.

2.1.7 Dry cooling tower. The glycol moves through the finned coil liquid-to-air heat exchanger, and is cooled by the colder outside air. When natural convection does not provide enough airflow, air is pulled through the heat exchanger fins by a large radial fan above the unit.

2.1.8 Batteries. The electrical storage consists of 40 absorbant glass mat (AGM) thin plate lead carbon batteries, of type NSB 100FT BLUE+ produced by Northstar, which hold a theoretical 30 kWh of electrical energy. The charging and discharging of these batteries is controlled by the MGS microgrid inverter.

2.1.9 Measurement and control system. This rack cabinet houses the servers and network components responsible for collecting sensor data and sending it downstream to storage and control functions. It also houses the power distribution equipment which divides incoming 3-phase power from the MGS into separate circuits enabling overcurrent protection and power measurement of individual system components.

2.1.10 Solar panels. The microgrid setup includes an array of 40 Eagle MX (JK07B) solar panels on the roof of the facility, produced by Jinko Solar, and use maximum power point tracking (MPPT) technology to maximize power extraction. They are designed to generate 265 Watts per panel with a maximum of 10 kW of solar power. The power is supplied to the microgrid solar inverter to be consumed, stored in batteries, or potentially sold to the utility grid.

2.1.11 Microgrid system. At the heart of the micro grid sits an ABB MicroGrid Solution (MGS100-40/27.6) controller unit, which regulates the flow of electricity in the system. It connects the system with the grid, solar panels and batteries. This unit allows the seamless transition from operating using the grid power to running from the batteries in the event of a grid failure. The unit is rated for a maximum load of 40 kW, a maximum of 27.6 kW photovoltaics output power, battery charging with 24 kW and a battery capacity of 276 kWh.

2.2 Dimensioning

The EDGE lab setup is dimensioned to resemble what is believed will be a realistic size of a future common edge data center with 10 kW of installed IT power. To make it possible to run the data center off-grid during a potential 3-hour peak period, a battery stack of 30 kWh was chosen. To prolong the potential period of off-grid operation, and reduce the grid power needed, 10 kW of photovoltaic panels were added to the microgrid.

To reduce the need for operating active cooling equipment when on battery or solar power and hence even further prolong off-grid operation, a 2 m³ insulated thermal energy storage (TES) tank was added to the system. It is installed between the chiller and the data center module and enables the use of stored cold water while e.g. in off-grid mode.

A simple calculation example assuming the IT running at full power for 3 hours (30 kWh) demonstrates how to determine tank temperature set-point from the relation $Q = mc_p\Delta T$: first insert the tank mass of water (2000kg) and its heat capacity (4.186kJ/kgK)

and then solve for $\Delta T = 12.9^\circ\text{C}$. A temperature differential of 6°C between the incoming water and cold aisle setpoint temperature is used, therefore when the cold aisle setpoint is 23°C , the storage tank should be at 4°C to be usable for 3 h off-grid.

The chiller, which includes both a compressor-based cooler and a free cooling heat exchanger, is heavily over-dimensioned for the size of the data center, given the data-sheets of the chiller unit. The main reason for using this over-dimensioned chiller is that a 10kW chiller with free cooling functionality was not commercially available. However, as the chiller-system and the connected dry-cooler are over-dimensioned, it enables free cooling operation at a higher ambient temperature, which reduces the time the chiller is in compressor-based cooling.

2.3 Measurement collection

Data from the sensors in the hardware equipment or hosted natively by the chiller or MGS unit is sampled approximately every 30 seconds by a common data collection system.

IT measurements. Data collected from Dell PowerEdge R430 servers includes temperature measurements for the CPUs and other on-board components, fan speeds and PSU parameters. The PDUs measure current, voltage, power consumption and power factor on all input phases and all outlets to the IT equipment. Network traffic is also monitored, albeit not used for this study.

Chiller, tank and flows. The chiller native interfaces allow data communication over Modbus RTU to include measurements of outdoor air temperature, dry cooler fan speed, free cooling valve position, temperatures for primary (glycol) and secondary (water) inlet and outlet to chiller, free cooling water temperature, operating mode and setpoint for secondary outlet (to the TES tank).

Volumetric flow and heat transfer are directly measured by ultrasonic heat meters (Kamstrup Multical 403) in both the primary and secondary chilling loops. Eleven PT100 temperature sensors were positioned in thermowells at different heights to capture the effects of stratification in the TES tank.

Microgrid controller. The MGS collects many measurements that can be accessed over Modbus TCP; those relevant for control or analysis in this study are: 1) Real, reactive and apparent grid power usage, energy import/export towards grid, and power factor, 2) Solar power input to inverter and its output, 3) UPS status, remaining battery time, as well as battery temperature, voltage and current, and 4) Real, reactive and apparent power usage to the rest of the EDGE setup, as well as its power factor and energy consumption.

Facility power. An ABB CMS-700 digital power meter is used to measure the individual power consumption of the chiller, dry-cooler, module, measurement network and module pump. This measures voltage and current and calculates the power using an estimated power factor of 0.97.

Estimated, aggregated or calculated quantities. The power consumption of the CRAH fan is not directly measured, but instead estimated as the total module power consumption minus the IT-load power given by the PDUs. Power Usage Effectiveness (PUE) is an example of another calculated measure, for which time averages over different horizons are also provided.

Other facility and environment measurement. The outside temperature is measured in proximity to the cooling tower. The ambient temperature inside the facility is also measured by a PT100 sensor placed near the tank. A silicon irradiance sensor placed on the roof next to the PV cells provides additional data on their production.

Software components. The open source software Zabbix¹ collects data from the servers, PDUs, network equipment, environmental sensors, chiller unit, microgrid controller as it can conveniently communicate with sensors through the MODBUS TCP and SNMP protocols, either directly communicating with the device or through a MODBUS TCP gateway. On-board sensors provided on the servers can furthermore be monitored by installing operating system resident agents responsible for communication with Zabbix.

The measurement data is fed into a Kafka streaming platform before being passed into a time-series data base (KairosDB) for long term storage and access by the visualisation (Grafana) and analysis (Jupyter, Python or Matlab) tools. The data collection solution used in the lab is designed for scalability and described in previous work [7], which we refer to for details on the full software stack.

2.4 Control system

The EDGE control system consists of both custom controllers and the factory controllers in the chiller and the MGS unit. The custom controllers are physically located on the measurement and control server on-site (unit 9 in Figure 1) which can be remotely accessed using a web-based UI or an SSH client. The various set-points and PID parameters can be changed through the UI, where the most recent measurements of the controlled variables can also be seen. Larger changes to the control algorithms are made by reprogramming the logic through the SSH connection. The factory control logic regulates the components inside the chiller, as well as the dry cooler fans, according to the rules described in the operation manuals [3, 4].

2.4.1 Data center aisle differential pressure controller. The CRAH fan is used to control the differential pressure (cold aisle pressure minus the hot aisle pressure) between the two data center aisles using a PI controller for regulation. Higher differential pressures result in the air mixing better, the aisles having more uniform temperatures, and affect the magnitude and direction of heat leakages between the aisles and to the ambient air.

2.4.2 Cold aisle air temperature controller. The secondary user side water pump is used by a PI controller to control the cold aisle air temperature. By adjusting the water flow rate through the CRAH, the heat removal rate can be adjusted, which allows control of the discharge air temperature.

2.4.3 Shunt valve controllers. The shunt valves on the primary and secondary side are both controlled with PI controllers. The primary shunt valve is used to keep the glycol entering the chiller above 5 °C so as to not freeze the other components where water is used. The secondary shunt valve is used to increase the water temperature pumped to the CRAH, in case the storage tank contains colder water than is required for regulating the module air temperature.

¹<https://www.zabbix.com>

2.4.4 Chiller factory control. The PLC inside the chiller unit exposes many of its internal parameters, either through Modbus TCP or by a touch screen interface. An example of the former is the water output temperature setpoint, which is used to control the temperature of water input to the TES tank.

The difference between outdoor air temperature and user side water temperature determines the chiller operation mode. Free-cooling is used when the outdoor temperature is colder than the water by a specific margin. A small hysteresis is used for returning to chiller mode. Partial free cooling mode is used when the system is in free cooling mode, but requires activating the compressors to reach the water output setpoint temperature. In this case the 3-way valve splits the flow of glycol between the free cooling and condenser heat exchangers.

In chiller mode the water output temperature is regulated by turning the compressors on and off as needed. Due to the chiller being over-dimensioned for the IT load, only one of the two compressors is required at a time. In free cooling mode the temperature regulation is done by either limiting the glycol flow rate to the free cooling heat exchanger using the 3-way valve or increasing the glycol pump and drycooler fan speed. The chiller water pump holds a constant flow rate, whereas the glycol pump varies its flow rate.

In chiller mode the drycooler fan speed is regulated linearly depending on the condensing temperature of the refrigerant inside the chiller. In free cooling mode its speed is regulated based on the temperature difference of the incoming water temperature to the output temperature setpoint.

2.4.5 IT load control. The IT load is entirely synthetic, and controlled on each individual server and CPU core using open source stress testing tools (stress-ng and mprime).

2.4.6 Microgrid control. The MGS is controlled through the default ABB software which automatically switches to battery and solar power in the event of a grid power outage, but also allows manually switching to off-grid operation. As it is proprietary, there were limits to the opportunities for under-the-hood modifications and some controls had to be carried out via the touchscreen user interface.

3 RESULTS

This section presents the results from simple experiments carried out during the second half of March 2020, when outside day time temperatures at the test site mostly kept around 0 °C, allowing the chiller unit to operate mostly in free cooling mode, but also sometimes in mixed mode for days that saw above 5 °C temperatures.

3.1 Experiment descriptions

The off-grid capabilities were investigated by comparing the energy consumption, battery discharge and TES tank temperatures with different off-grid strategies. The 'baseline' for the comparisons is a period where the data center operates in the 'normal' operation conditions, where the setpoints are as listed below.

- IT load at 30 %
- Cold aisle temperature setpoint at 23 °C
- Differential pressure setpoint at 0 Pa
- Chiller output water temperature setpoint at 10 °C

Experiments were carried out to compare different strategies for utilizing the off-grid capabilities of the EDGE data center. In the experiments, measurement began 1 hour before the grid power to the EDGE system is cut off for 2.5 hours, after which the grid power is reconnected for 4.5 hours. All experiments are done at the same time of day to minimize variations in the surrounding temperature.

3.1.1 Experiment 1. The first experiment investigates how the system behaves when the grid power is disconnected and the system switches to battery power, but everything continues as usual.

While off-grid, the batteries are partially drained, and the solar panels also provide some power. Most of the power is used to keep the IT equipment running, but some is also used to cool the module and storage tank. When grid power is reconnected, the batteries are recharged slowly, which increases the power consumption slightly above the normal level.

3.1.2 Experiment 2. In comparison to the first experiment, in the second experiment the chiller is turned off when the grid power is cut off, so that the cold water stored in the TES tank may be consumed for cooling the module instead of using battery power to power the chiller.

During this experiment the batteries are drained slightly slower, as the cooling power consumption is lowered with only secondary side units active since the chiller and tower fan are put in idle mode. When the grid power is reconnected, the batteries are recharged, and the TES tank is recharged, which increases the power consumption compared to experiment 1.

3.1.3 Natural experiment: chiller failure. At one point during our preparations the chiller unit was switched off so that maintenance could be carried out (to refill water for the secondary loop); we left the IT load on as this constitutes a useful natural experiment that is presented in detail later in this article as an example of a failure scenario.

3.2 Results from experiment 1 and 2

3.2.1 TES Tank temperatures. As can be seen for TES tank temperatures from experiment 2 in Figure 4, initially the storage tank is mostly full of $\sim 10^\circ\text{C}$ water. After the chiller is powered off temperatures increase slowly from top to bottom, indicating a significant degree of stratification in the warming up of the water in the tank. The average temperature for the sensors in the tank increases from 11.4 to 19.5°C .

Figure 4 shows that after the chiller is switched back on, it takes up to 5 hours for the temperatures at all but the top level in the TES tank to return to their original levels. It is noted that most of heat removal work is done within the first hour of re-cooling.

3.2.2 Chiller power. The first column of Table 2 shows the average power consumption for chiller and tower fan over the experiments. For experiment 1 the chiller is in operation throughout and power consumption at any time is close to the average levels. This is not the case for experiment 2, where the chiller and cooling tower go idle from one hour and their consumption decreases to a minimum at approximately 210 W and 140 W , respectively (idle column in Table 2). After the chiller is started again at around the 3.5 hours mark, the chiller unit must work harder to expel the excess heat,

especially during the first hour for which the cooling power is markedly increased to above 1600 W and 1300 W , respectively for the chiller and tower (recover column in table). This observation is consistent with the hour-long phase of accelerated cooling noted in the last section for the tank temperatures.

Figure 4: In experiment 2, the TES tank temperatures (solid) rise during the black-out since its water is used for cooling the DC module. Mean cold aisle temperature (dashed) stays constant.

Table 1: Mean chiller power consumption by period.

	All (W)	Idle (W)	Recover (W)
Experiment 1: chiller	1045	-	-
tower	172	-	-
Experiment 2: chiller	895	214	1661
tower	342	138	1350

The total energy used for cooling over experiment 2 increases slightly (by 0.5 kWh 24 hours after experiment start) when the cooling equipment is turned off and back on, compared to when they continue working through the simulated power failure.

3.2.3 Battery draw. The available sunshine varied for the different days the experiments were carried out. For fair comparison, a hypothetical *nightly* scenario is therefore also presented, which subtracts all PV production from the actual battery use. Table 2 shows both actual and *nightly* (solar compensated) power consumption for the experiments. The *nightly* draw from the batteries thus drops from about 9.8 kW (experiment 1) to 8.9 kW (experiment 2), that is about 900 W , which translates into a total energy saved of about 2.25 kWh by cooling from the tank during the simulated 2.5 hour black-out.

3.2.4 Heat rejection. Estimates for the heat rejected in the system displayed in Figure 5 show that the direct heat rejection (yellow curves) from IT use is the same for the three experiments as the module's CRAH is still provided by cold water from the tank. Cooling for the tank (blue curves) is halted during the black-out showing

Table 2: Battery draw (* nightly in last column).

	Total draw (kWh)	Solar input (kWh)	Power (kW)	Power* (kW)
Experiment 1	22.1	2.47	8.83	9.82
Experiment 2	16.1	6.09	6.43	8.86

Figure 5: Estimates for cumulative heat rejected from IT (yellow) and tank (blue) for experiment 1 (solid) and 2 (dashed).

a flat curve for experiments 2 (dashed in figure), after which heat rejection is accelerated to bring back the tank to the setpoint. Note a small (accumulating) difference between the IT use and the heat rejection from the tank that is likely due to additional heat losses to the environment.

3.3 Natural experiment: chiller failure

The results from the natural experiment are displayed in Figure 6, which shows the progress of temperatures measured in the TES tank after the chiller unit is switched off at approximately 30 minutes into the displayed episode; at around 3 hours a noticeable change in the cooling performance when the water in the tank has been circulated once to the edge module. Note that there is a small dip in the interval 3–3.5 hours due to the chiller unit being briefly switched back on for testing.

After this 3 hour mark the temperatures start rising quickly in the module: the experiment showed that the hottest spot at the top of the racks reaches above 32 °C, which is above the ASHRAE recommended range for data centers [1, 2]. This occurs about 5 hours into the experiment when also the top hot aisle temperatures have reached high levels - approximately 48 °C - here the experiment was aborted by opening the doors to the DC module.

Cold air from the surrounding laboratory room entered the DC, causing the module to cool rapidly and halting the increase in tank temperatures (6). The chiller unit is switched on again at around 7.5 hours after which the water in the TES tank is quickly cooled again.

3.4 Analysis

Tank temperatures. The storage tank temperature change, ΔT , is calculated as the rise in the average temperature from the time

Figure 6: For a natural experiment when chiller failed, the tank temperatures (solid) rise gradually in two identifiable phases: stratified and (almost uniformly) mixed. The average cold aisle temperature (dashed) rise quickly after the cold tank mixes (around 3 hours).

the setup goes off-grid and its value 2.5 hours later. For experiment 1 the difference is zero, as the storage tank is not used, while for experiment 2 the ΔT is 8.1 °C. An upper limit for the average temperature change can be estimated at 8.7 °C by taking the total heat injected by the IT equipment in the module over the total thermal capacity of the water in the tank (i.e. heat capacity of water times the water mass). We note the observed increases in the average tank temperature are below this bound.

Off-grid battery use time. It was noted earlier that the average power drawn from the battery when the system is off-grid is about 900 W lower for experiment 2 than for experiment 1. This points to the benefit of switching the chiller off during blackout and use of the TES tank to increase the potential off-grid operation time. By extrapolating the linear trend of the battery discharge until the battery charge reaches zero to give a maximum off-grid running times of 3.1 hours for experiment 1 and 3.4 hours for experiment 2.

Thus for a system of similar dimensions as the present set-up with 30 kWh battery, the increase in off-grid duration is about 10%, or a modest 20 minutes. This estimate is however based on similar conditions as in the experiments - 30% IT load, cool outdoor temperature and free cooling - one would expect that the difference in off-grid operation would be larger for days when the chiller must operate with compressor-based cooling.

Observations of the system under the same load, but when environmental conditions required compressor cooling, showed that typical power consumption by the chiller unit to be around 2.2 kW, which is about twice that for free cooling mode. In this case the tank extends the off-grid operation time by about 40 minutes.

Off-grid tank warm-up time. From the natural experiment displayed in Figure 6 it was observed that (at this load) the tank can last for a period of approximately 5 hours, which roughly separates into two 2.5 hour regimes; in the first phase the data center temperatures are kept near the normal operation setpoint, while

in the second they gradually drift higher until they are above the recommended ASHRAE safe limit.

If the cold aisle temperature limitation of 27 °C is lifted even higher, for example to the server manufacturers specification upper limit of 35 °C [5], the storage tank may be used to cool the data center for longer. In theory, the servers can be used until they are turned off when the CPU temperature reaches 90 °C.

Overhead of using energy storage. Lead acid batteries suffer from self-discharge at a typical rate of 3-5% per month and must be regularly topped up; charging by itself only has an approximate efficiency of 85%. The electrical energy consumed to keep the batteries topped up is therefore roughly 1.5 W.

The rate of loss cooling capacity from the TES tank can be estimated from its surface area, the thermal conductivity of the insulated wall and the temperature difference to the environment. Assuming it is kept about 10 °C colder than the room the thermal losses are on the order of 100 W. Converting the thermal losses to electrical losses, roughly 11W of electrical power is required to maintain the thermal losses using free cooling.

Space effectiveness. The TES tank stores 15 kWh/m³ of cooling capacity assuming the ΔT of the tank rises 13°C, compared to the 23 kWh/m³ of electrical power the batteries contain. The amount of electrical energy that the thermal energy corresponds to depends on the efficiency of the cooling equipment, which further depends on the outside temperature. However, the capacity of the TES tank is not fixed, and can be increased by increasing the ΔT of the stored water and cold aisle temperature.

4 DISCUSSION

The overhead of both TES tank and batteries decreases the overall energy efficiency of the entire system, however these losses have to be weighed against the reliance benefits they provide in allowing the system to continue running through a blackout. When the overhead is given a monetary value, it can be seen that these systems are cheap to operate. The increase to the operational costs from storing energy to the batteries and TES tank amounts to an increase of ~11 €/year, assuming an electricity cost of 0.1 €/kWh. This is a small price to pay for allowing the system to continue operation for 3 hours into potential power outages, subject to capital expenditure assessment. It is less than 0.2% if set in relation to IT power consumption and corresponds to a small increase in the PUE.

How to strike the balance between energy efficiency and autonomy ultimately depends on the specific use case. Where the battery pack must be sized according to the desired off-grid operating horizon, a TES tank can be used to extend the horizon by dimensioning it to provide cooling for as long as the batteries allow the IT to run.

From the results presented earlier, it is clear that the battery storage is more energy efficient to maintain at full capacity, more space efficient in terms of kWh stored per m³, and more flexible as it can be used for both IT and cooling. However the TES tank offers opportunities for planned cooling such as overnight free-cooling storage to be used in the at day to avoid running the compressor. A variable electricity price in addition to temperature difference between night and day would further increase the potential for cost savings, and could motivate the investment in TES tanks.

5 CONCLUSIONS

This article presents the EDGE lab for small microgrid data centers with solar panels, batteries and coolant thermal energy storage tanks. The setup is described in detail on hardware for both cooling, power transmission and IT, as well as software used for monitoring, control and data collection.

The arrangement is used as a testbed for exploring the interplay between IT load, renewable power, battery storage and thermal energy storage, which is demonstrated in an experiment that compares different strategies for off-grid operation by planned use of battery and TES tanks. Different variables important for the right-sizing of batteries and TES tank are presented supported by analysis of a set of experiments.

In conclusion, a setup that includes thermal energy storage with a microgrid connected to an edge data center can contribute to its resilience by adding operational hours of the IT during a blackout or planned maintenance.

Future work. The comprehensive data collection from all parts of the arrangement set the stage for detailed modelling, which has been verified for the thermal loops in the initial study [8] and is an area for further exploration.

The use of the TES tank as a means to actively manage costs by moving cooling work to night time when outside temperatures and electricity costs are lower is worth examining in more detail.

The microgrid setup with PV-cells and batteries also offers opportunities for future modelling and experimental work on the planned use of electrical energy and cooling with forecasts of IT load, weather and future electricity prices.

ACKNOWLEDGMENTS

This study was supported by Vinnova grant ITEA3-17002 (AutoDC), and the Swedish Energy Agency grant 2016-007959 (DMI/SampEL). The authors also thank the following companies for their generous support in building the testbed: Box Modul AB, Bensby Rostfria AB, Borö Pannan AB, Enoc System AB, CEJN AB and ABB Ltd.

REFERENCES

- [1] ASHRAE Technical Committee (TC) 9.9 Mission Critical Facilities, Data Centers, Technology Spaces, and Electronic Equipment. 2016. *Data Center Power Equipment Thermal Guidelines and Best Practices*. Whitepaper. ASHRAE.
- [2] ASHRAE Technical Committee (TC) 9.9 Mission Critical Facilities, Technology Spaces, and Electronic Equipment. 2011. *2011 Thermal Guidelines for Data Processing Environments – Expanded Data Center Classes and Usage Guidance*. Whitepaper. ASHRAE.
- [3] Blue Box Group S.r.l. 2017. *Technical catalogue - Tetris W Rev FC/NG*. Blue Box Group S.r.l.
- [4] Blue Box Group S.r.l. 2018. *Controller manual - Service, Series: Zeta Rev Series, Beta Rev Series, Tetris 2 Series, Tetris Rev W LC, Tetris Rev W LC/HP*. Blue Box Group S.r.l.
- [5] Dell Inc. 2018. *Dell PowerEdge R430 Owner's Manual*. Dell Inc.
- [6] Íñigo Goiri, William Katsak, Kien Le, Thu D. Nguyen, and Ricardo Bianchini. 2013. Parasol and GreenSwitch. In *Proceedings of the eighteenth international conference on Architectural support for programming languages and operating systems - ASPLOS '13*. ACM Press.
- [7] Jonas Gustafsson, Sebastian Fredriksson, Magnus Nilsson-Mäki, Daniel Olsson, Jeffrey Sarkinen, Henrik Niska, Nicolas Seyvet, Tor Minde, and Jonathan Summers. 2018. A demonstration of monitoring and measuring data centers for energy efficiency using opensource tools. In *Proceedings of the Ninth International Conference on Future Energy Systems (Karlsruhe, Germany) (e-Energy '18)*. Association for Computing Machinery, New York, NY, USA, 506–512.
- [8] Mikko Siltala. 2020. *Simulating data center cooling systems: data-driven and physical modeling methods*. Master's Thesis. Aalto University.