Article

Data Centers Optimized Integration with Multi-Energy Grids: Test Cases and Results in Operational Environment

Tudor Cioara [1], Marcel Antal [1,*], Claudia Daniela Antal (Pop) [1], Ionut Anghel [1], Massimo Bertoncini [2], Diego Arnone [2], Marilena Lazzaro [2], Marzia Mammina [2], Terpsichori-Helen Velivassaki [3], Artemis Voulkidis [4], Yoann Ricordel [5], Nicolas Sainthérant [5], Ariel Oleksiak [6] and Wojciech Piatek [6]

[1] Computer Science Department, Technical University of Cluj-Napoca, Memorandumului 28, 400114 Cluj-Napoca, Romania; tudor.cioara@cs.utcluj.ro (T.C.); claudia.pop@cs.utcluj.ro (C.D.A.); ionut.anghel@cs.utcluj.ro (I.A.)

[2] Engineering Ingegneria Informatica, Piazzale dell'Agricoltura 24, 00144 Rome, Italy; massimo.bertoncini@eng.it (M.B.); diego.arnone@eng.it (D.A.); marilena.lazzaro@eng.it (M.L.); marzia.mammina@eng.it (M.M.)

[3] Singular Logic, Achaias 3 & Trizinias st., Kifissia, P.C. 145 64 Attica, Greece; tvelivassaki@ep.singularlogic.eu

[4] PowerOps, Faraday Rd, Swindon SN3 5HQ, UK; artemis@power-ops.com

[5] Qarnot Computing, 40–42 Rue Barbès, 92120 Montrouge, France; yoann.ricordel@qarnot-computing.com (Y.R.); nicolas.saintherant@qarnot-computing.com (N.S.)

[6] Poznan Supercomputing and Networking Center, Jana Pawła II 10, 61-139 Poznan, Poland; ariel@man.poznan.pl (A.O.); piatek@man.poznan.pl (W.P.)

* Correspondence: marcel.antal@cs.utcluj.ro; Tel.: +40-264-202-352

check for
updates

Abstract: In this paper, we address the management of Data Centers (DCs) by considering their optimal integration with the electrical, thermal, and IT (Information Technology) networks helping them to meet sustainability objectives and gain primary energy savings. Innovative scenarios are defined for exploiting the DCs electrical, thermal, and workload flexibility as a commodity and Information and Communication Technologies (ICT) are proposed and used as enablers for the scenarios' implementation. The technology and scenarios were evaluated in the context of two operational DCs: a micro DC in Poznan which has on-site renewable sources and a DC in Point Saint Martin. The test cases' results validate the possibility of using renewable energy sources (RES) for exploiting DCs' energy flexibility and the potential of combining IT load migration with the availability of RES to increase the amount of energy flexibility by finding a trade-off between the flexibility level, IT load Quality of Service (QoS), and the RES production level. Moreover, the experiments conducted show that the DCs can successfully adapt their thermal energy profile for heat re-use as well as the combined electrical and thermal energy profiles to match specific flexibility requests.

Keywords: Data Centers; energy efficiency; heat re-use; energy flexibility; workload migration; evaluation in relevant environment

1. Introduction

As the ICT services industry is blooming, such services being requested in almost every domain or activity, Data Centers (DCs) are constructed and operated to supply the continuous demand of computing resources with high availability. However, this is the nice side of the story, because as this also has an environmental impact, the DCs sector is estimated to consume 1.4% of global electricity [1].

Thus, it is no longer sufficient to address the DCs' energy efficiency problems from the perspective of decreasing their energy consumption, but new research efforts are aiming to increase the share of renewable energy used for their operation and to manage them for optimal integration with local multi-energy grids [2,3].

Firstly, the DCs are large generators of residual heat which can be recovered and reused in nearby heat grids and offer them a new revenue stream [4]. This is rather challenging due to the continuous hardware upgrades that increase the power density of the chips, leading to an even higher energy demand for the cooling system to eliminate the heat produced by the Information Technology (IT) servers to execute the workload [5]. In addition to the potential hot sports that may endanger the safe equipment operation, other challenges are created by the relatively low temperatures of recovered heat compared with the ones needed to heat up a building and the difficulty of transporting heat over long distances [6,7]. Current studies show that the DCs' can offer a secured supply of heat accounting more than 60 TW h in Europe [8]. If the heat is not constantly dissipated, the temperature in the DC can exceed the normal operating setpoints and the IT equipment can be damaged. Thus, the cooling system is a significant contributor to the DCs' high energy consumption and even in well designed DCs, it takes almost 37% of the total energy consumption [9]. Several aspects are of most interest for DCs when providing waste heat to district grids [10]: improving the DCs' energy efficiency, better integration of renewable energy and downstream waste heat monetization. The adoption of new emerging technologies, such as machine learning or blockchain, force DCs to investigate even other possibilities for more efficient cooling and heat reuse, such as adopting liquid cooling [11]. This trend was accelerated with the introduction of AI-friendly processors, which consume lots of energy and their heat dissipation could not be managed any more using air cooling. Moreover, hybrid solutions featuring liquid cooling are adopted while at the same time re-using the heat in the smart energy grid [12]. These solutions require complex ICT-based modelling and optimization techniques to assess the DC potential thermal flexibility and associated operation optimization [13].

Secondly, the DCs are characterized by flexible energy loads that may be used for assuring a better integration with local smart energy grids by participating in DR programs and delivering ancillary services [14,15]. In this way, the DCs will contribute to the continuity and security of energy supply at affordable costs and grid resilience [16]. Moreover, if the renewable energy is not self-consumed locally, problems such as overvoltage or electric equipment damage may appear at a local micro-grid level and could be escaladed to higher management levels [17]. To be truly integrated in the grid supply, renewable energy sources (RES), since they are volatile, need of flexibility options and DR programs to be put in place [18]. Few approaches are addressing the exploitation of DCs' electrical energy flexibility to achieve better integration into the local energy grids [19]. Modern ICT infrastructures allow the development of demand and energy management solutions that will allow DCs to be involved scenarios such as reduction of peak power demand during DR periods [20]. DCs' strategies for providing demand response are usually referring to shutting down IT equipment, using Dynamic Voltage Frequency Scaling (DVFS), load shifting or queuing IT workload, temperature set point adjustment, load migration and IT equipment load reduction [21–23]. To increase their flexibility potential, the DCs may rely on non-electrical cooling devices such thermal storages for pre-cooling or post-cooling [6] and the DC IT workload migration in federated DCs [7]. Shifting flexibility to meet the demand by leveraging on workload scheduling usually involves the live migration of virtual machines, even in other DCs [24], or postponing delay-tolerant workloads to future execution points [25]. The migration of load to distributed partners' DCs as part of an optimization process is currently addressed through more promising techniques such as blockchain-based workload scheduling by trying to solve the data transfer security issues [26]. This typically involves different Artificial Intelligence (AI) algorithms and techniques with a prevalence of prediction heuristics [27].

Table 1 classifies the above-presented approaches in terms of proposed solutions for DC integration in smart grids and innovative techniques and technologies developed.

Table 1. Literature approaches comparison.

DC Multi Energy Grid Integration	State of the Art Techniques
Heat re-use	• Heat reuse models based on Computational Fluid Dynamics (CFD) [2,6,7] • Neural networks-based prediction techniques [6–8] • Optimization heuristics [2,12]
Electrical energy shifting	• Systems of Systems (SoS) modelling and simulation of DCs and their building components [3,13] • Optimization heuristics [3,13,14] • Electronic marketplace for trading flexible energy [2,14] • Energy consumption/production prediction techniques [15,20]
DC Federation	• IT workload execution time shifting [2,21,25] • Workload spatial relocation in federated DCs [2,24] • Optimization heuristics [21,24,25] • Blockchain-based DC workload distribution and management [26]

In this context, the innovative vision defined by H2020 project CATALYST [28] is that the DCs' energy efficiency, should be addressed by managing their operation considering the optimal integration with electrical, thermal and data networks. On the one hand, the DCs have good, yet mostly unexploited, potential regarding their energy demand flexibility, which makes them great potential contributors to the ongoing efforts for more efficient and integrated management of the smart grid. On the other hand, they are large producers of residual heat which can be recovered and re-used near district heating infrastructures. At the same time, they have great IT data network connections which may provide a new source of flexibility and optimization: the workload relocation to/from other DCs to meet some green objectives such as following the renewable energy. As shown above, the state-of-the-art approaches to DC energy efficiency address only partially the above-presented aspects lacking a holistic approach. The work presented in this paper contributes to the process of creating the necessary technological infrastructure for establishing active integrative links among DCs and electrical, heat, and IT networks which are currently missing.

In summary, the paper brings the following contributions:

- Defines innovative scenarios for DCs, allowing them to exploit their electrical, thermal and networks connections for trading flexibility as a commodity, aiming to gain primary energy savings and contribute to the local grid sustainability.
- Describes an architecture and innovative ICT technologies that act as a facilitator for the implementation of the defined scenarios.
- Presents electrical energy, thermal energy and IT load migration flexibility results in two pilot data centers, showing the feasibility and improvements brought by the proposed ICT technologies in some of the new scenarios.

The rest of the paper is structured as follows: Section 2 presents the new scenarios and ICT technology, Section 3 describes the results obtained in two pilot DCs, and Section 4 concludes the paper.

2. Scenarios and Technology

Table 2 shows the defined scenarios for efficient management and operation of the DCs at the crossroads of data, electrical energy and thermal energy networks trading electricity, energy flexibility, heat and IT load as commodities. The definition of scenarios is done incrementally: Scenarios 1, 2 and 3 considering the network connections in isolations, Scenarios 4, 5 and 6 consider combinations of the two network connections together, while Scenario 7 is the most complex one in which all three network connections are considered at once. The scenarios highlight the importance of various commodities

in obtaining primary energy savings and decreasing carbon footprint and the network connections providing the necessary infrastructure in achieving this.

Table 2. New scenarios defined.

Scenario	Optimization Objectives	Utility Network
Scenario 1: Single DC Providing Electrical Energy Flexibility	Optimize the DC operation to deliver energy flexibility services to the surrounding electrical energy grids ecosystems aiming to create new income source and reduce DC energy costs. Assess resiliency of energy supply and flexibility, against adverse climatic events or abnormal demand, trading off DC assets energy generation or consumption against on-site or distributed RES, energy storage and efficiency.	Electrical Energy Network
Scenario 2: Single DC Providing Heat	Optimize DC operations to deliver heat to the local heat grid. Recover, redistribute and reuse DC residual heat for building space heating (residential and non-residential such hospitals, hotels, greenhouses and pools), service hot water and industrial processes. The DC achieves significant energy & cost savings, reduces its CO_2 emissions, contributes to reducing the system-level environmental footprint and supports smart city urbanization.	Thermal Energy Network
Scenario 3: Workload Federated DCs	Exploit migration of traceable ICT-load between federated DCs, to match the IT load demands with time-varying on-site RES availability (including Utility/non-Utility owned legacy assets) thus reducing the operational costs and increasing the share of renewable energy used.	IT Data Network
Scenario 4: Single DC Providing both Electrical Energy Flexibility and Heat	Optimize the DC operation to deliver both electrical energy flexibility services and heat to the surrounding energy (power and heat) grids ecosystems. The DC will act as convertor between electrical and thermal energy and vice versa to gain extra revenue on top of normal operation.	Electrical Energy & Thermal Energy Networks
Scenario 5: Workload Federated DCs Providing Electrical Energy Flexibility	Exploit migration of traceable ICT-load between federated DCs to deliver energy flexibility services to the surrounding power grids ecosystems aiming to increase DC income for trading flexibility.	Electrical Energy & IT Data Networks
Scenario 6: Workload Federated DCs Providing Heat	Exploit migration of traceable ICT-load between federated DCs to deliver heat to their local heat grids aiming to increase the revenue for the reuse of their residual heat.	Thermal Energy & IT Data Networks
Scenario 7: Workload Federated DCs Providing Both Thermal and Electrical Energy Flexibility	Exploit migration of traceable ICT-load between federated DCs to deliver: (i) heat to the surrounding thermal grids and (ii) energy flexibility to the surrounding power grids.	Electrical Energy & Thermal Energy & IT Data Networks

To address this innovative vision, several ICT technologies have been developed and coherently integrated in a framework architecture consisting of three interacting systems (see Figure 1): the DC Flexibility Manager, DCs Federation Manager, Data center infrastructure management (DCIM) and Utility Networks Integration.
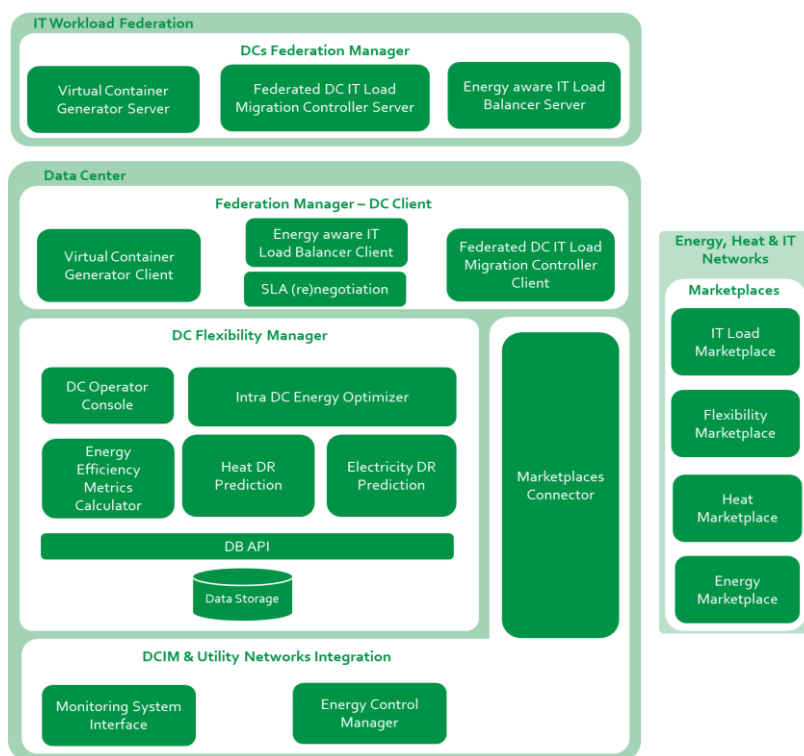
Figure 1. Data Centers' (DCs) optimization framework architecture.

2.1. DC Flexibility Manager System

The DC Flexibility Manager sub-system is responsible for improving DC energy awareness and energy efficiency, by exploiting the internal flexibility for providing energy services to power and heat energy grids. The main components of this sub-system are detailed in Table 3.

Table 3. DC Flexibility Manager components.

Component	Objective	Relevant Techniques/Technologies
Intra DC Energy Optimizer	Decides on the optimization action plans that will allow the DC to exploit its latent energy flexibility to provide electricity flexibility services in its micro-grid, to re-use heat in nearby neighborhoods, and finally to leverage on workload reallocation in other DCs as potential source of additional energy flexibility	Optimization heuristics, DC model simulation
Electricity DR Prediction	Predicts the DC energy consumption, generation and flexibility using the following time windows: day ahead (24 h ahead), intraday (4 h ahead) and near real time (1 h ahead).	Machine learning based models
Heat DR Prediction	Predicts heat available to be re-used in nearby neighborhoods. Estimates the temperature of the hot air recovered by the heat pumps from the server room in various configurations and use the data to train Multi-Layer Perceptron prediction model.	Computational Fluid Dynamics, Neural Networks
Efficiency Metrics Calculator	Calculates different metrics in close relation with decided optimization plans aiming to assess their impact onto the DC operation. Example of metrics used: Adaptability Power Curve at RES, Data Centre Adapt, Grid Utilization Factor, Energy Reuse Factor, etc.	
Data Storage	Stores the DC main sub-system characteristics, thermal and energy monitored data, prediction outcomes and optimization action plans.	NoSQL database
DC Operator Console	Displays information on the monitoring, forecasting as well as on flexibility optimization decision making. It allows the DC operator to select and validate an optimization plan and to configure the DC optimization strategies.	React JavaScript library

The interaction between the DC Flexibility Manager system components is presented in Figure 2. The Intra DC Energy Optimizer component takes as input the thermal and electrical energy predictions determined by the Electricity and Heat DR Prediction components as well as the DC model describing the characteristics and operation. Its main output is the optimal flexibility shifting action plan, whereby the DC energy profile is adapted to provide various services in the Electrical and Heat Marketplaces. The Energy Efficiency Metrics Calculator estimates the values of the metrics and the optimization action plans are displayed on the DC Operator Console for validation. If the plan is validated by the operator, its actions will be executed.
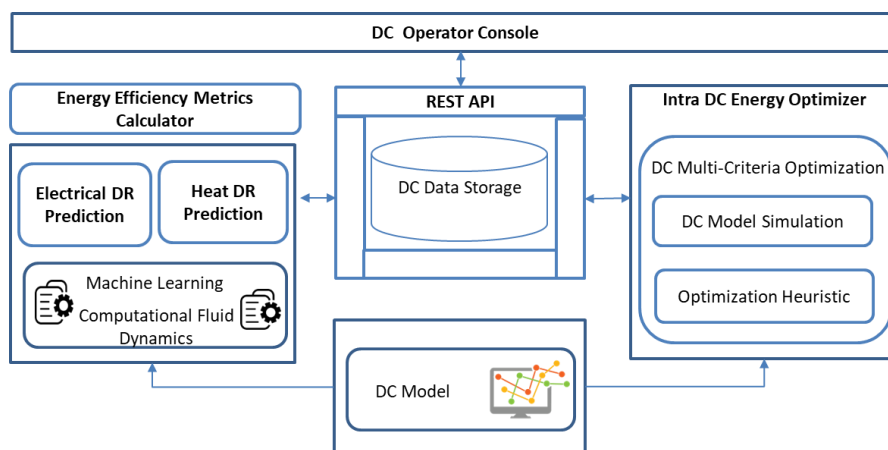

Figure 2. DC Flexibility Manager System architecture.

2.2. DCs Federation Manager

DCs Federation Manager is responsible for orchestrating the workload relocation among federated DCs. Table 4 describes the main components of this system.

Table 4. DCs Federation Manager components.

Component	Objective	Relevant Techniques
Energy Aware IT Load Balancer	Decides on the optimal IT loads placement across the federation in DCs offering capacity in the most efficient or green way. It defines the actual IT loads relocation plan based on the capacity offers and bids of other DCs to ensure the implementation of follow the renewable energy or minimum energy price strategies.	Knapsack algorithm, Branch and Bound techniques
Federated DC IT Load Migration Controller	Performs the actual live IT load migration between federated DCs, which belong to different administrative domains ensuring almost zero downtime.	Live migration of VMs
Virtual Container Generator	A distributed (reversed client-server model) component responsible for tracking information related to the lifetime of IT virtual loads, that is virtual machines or containers, on the blockchain, effectively transforming them into virtual containers (VCs).	Blockchain for IT load traceability
SLA (re)negotiation	Responsible for monitoring the SLA compliance of the CATALYST VCs, based on the events registered in the VCG. SLARC operation is based on the notion of Service Level Objectives (SLOs) referring to levelled acceptable behavior against a target objective for given periods. SLARC would notify registered parties prior to SLA breakage.	Blockchain, Publish/Subscribe mechanisms

The Federated DC MIgration Controller (DCMC) [29], including Master and Lite Client and Server, enables the live migration of IT load among different administrative domains, without affecting the end users' accounting and without service interruption (see Figure 3). The secure communication channels between DCs are deployed by the DCMC through OpenVPN and secured by Oauth2.0 to

ensure secure and authorized transfers. A new VM in the form of a virtual compute node gets created in an OpenStack environment of the destination DC to host the migrated load at the destination, while tokens for authenticating the DCMC components are retrieved via an integrated KeyCloak server (part of DCMC software). The DCMC software can attach the new virtual compute to the source DC's OpenStack installation, which means that the source DC is the sole owner/manager of the load even albeit in a different administrative domain.

Figure 3. DCMC architecture.

In detail, the DC IT Load Migration Controller clients will receive migration offers or bids requests by the Energy-aware IT Load Balancer, originally spurred by the Intra DC Energy Optimizer. The DC IT Load Migration Controller client will also inform the Energy Aware Load Balancer of potential rejection or acceptance. Upon acceptance of the migration bid/offer, DC IT Load Migration Controller clients will inform the DC IT Load Migration Controller server that they are ready for migration; in case of a migration bid, the DC IT Load Migration Controller client will first set up the virtual compute node to host the migrated load. Then, DC IT Load Migration Controller server is responsible for setting up a secure communication channel between the two DCs: the source and the destination of the IT load. After communication is set up and the source DC is connected to the virtual compute node at the destination DC, the migration starts. Then, the DC IT Load Migration Controller clients inform both the Energy Aware IT Load Balancer about the success or failure of the migration.

Under a federated DCs perspective, the Virtual Container Generator enables IT loads trackability and allows for indisputable SLA monitoring. In short, the Virtual Container Generator client will offer information related to the lifetime of the VCs, which will first be translated by the SLA (re)negotiation component into service license objective (SLO) events and second into SLA status compliance. This information will be fed to the Energy Aware IT Load Balancer component so that effective decisions on Virtual Container load migrations may be achieved. The SLA (re)negotiation component features two methods of information retrieval, which are RESTful and publish-subscribe, while for data feeding, pull requests towards the Virtual Container Generator will be performed.

2.3. DCIM and Utility Networks Integration

This system offers different components for allowing integration with existing Data center infrastructure management software or with the utilities networks considered for flexibility exploitation: electricity, heat and IT.

The Marketplace Connector acts as a mediator between the DC and the potential Marketplaces that are setup and running and on which the DC may participate. Through the Marketplace Connector, a DC can provide flexibility services, trade electrical or thermal energy and IT workload. Additionally, it provides an interface, and the electricity or energy grids operators are responsible for listening after Demand Response (DR) signals sent by the Distribution System Operator for reducing the DC energy consumption at critical times or for providing heat to the District Heating network. The DR request is

forwarded to the Intra DC Energy Optimizer, which evaluates the possibility for the DC to opt in the request based on the optimization criteria.

Energy Control Manager interfaces the DC appliances and local RES via existing DC infrastructure management system (DCIM), or other control systems (e.g., OPC server), implements, and executes the optimization action plans.

The Monitoring System Interface interacts with existing monitoring systems already installed in a DC, adapts the monitoring data received periodically, and provides this data to the Data Storage component from which it will be analyzed in the flexibility optimization processes.

3. DC Pilots and Results

In this section, we show the validation results of the proposed technology throughout the first four scenarios defined in Table 2 in two pilot DCs, a Micro Cloud DC connected to PV system and located in Poznan Poland and a Colocation DC located in Pont Saint Martin, Italy. The selection of the scenarios to be evaluated was driven by the hardware characteristics of the considered pilots DCs, their type and sources of flexibility available. For the real experiments, relevant measurements were taken from the pilot DCs, flexibility actions were computed by the software stack deployed in the pilot and the actions were executed by leveraging on the integration with existing DCIM (see Table 5). Finally, relevant KPIs were computed to determine the energy savings and thermal and electrical energy flexibility committed. The KPIs were calculated exclusively, considering the monitored data reported in the experiment and no financial or energetic parameters were assumed.
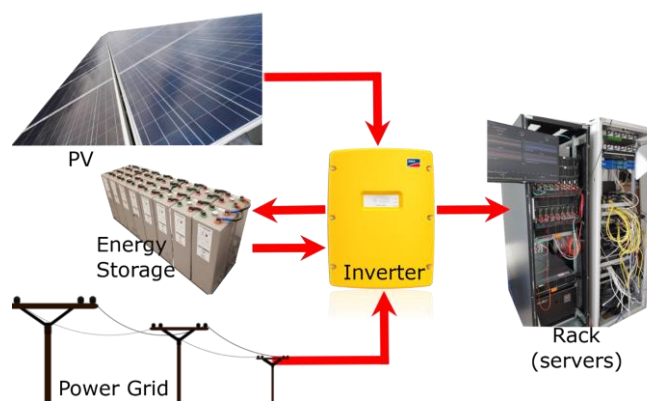
Table 5. Scenarios mapping on DC pilots.

Scenario Number	Poznan Cloud DC	Pont Saint Martin Colocation DC
1	🟩	🟦
2		🟩
3	🟩	
4		🟩
5		
6	🟦	
7		

Legend: Green: Real experiments described in next sub-sections; the flexibility actions are executed in the pilot DCs. Blue: implicitly covered; not further detailed.

3.1. Poznan Micro DC with Photovoltaic System

The configuration of the DC pilot test bed is presented in Figure 4. It features two racks with 50 server nodes which consume approximately 8.5 kW of power in a maximum load and about 3 kW in idle state. The servers' utilization varies only a few regular services running on the servers; thus, most of the workload (around 90%) constitutes tasks that can be easily shifted.

Figure 4. Pilot site: micro DC connected to the photovoltaic system with energy storage.

The racks are connected to the photovoltaic system, which consists of 80 PV panels with 20 kW peak power. Servers are directly connected to inverters so that they can be supplied: directly by power grid, using renewable in the case of enough generation or by batteries (75 kWh) if energy production is too low. If batteries are discharged below 60%, servers immediately switch to power grid to extend battery life. The maximum load of the PV used by servers, due to electrical constraints, is limited to 7.3 kW. With fully charged batteries and no solar radiation, the battery could keep the completely loaded servers for 4 h.

The cooling system consumes up to 3.2 kW with a cooling capacity of 10 kW (maximum 11 kW). The testbed DC system characteristics are summarized in Table 6.

Table 6. DC system characteristics.

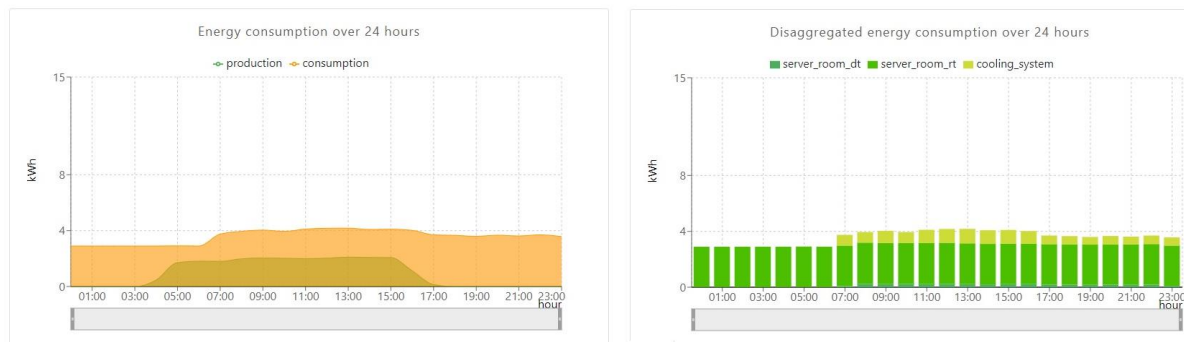| Sub-System	Characteristics
IT Servers	**1st rack: 26 nodes** (18 Xeon E5, 4 Taishan, 4 Tesla K80), energy consumption between ~6.5 kW max and ~2.5 kW on idle **2nd rack: ~32 nodes** (6x Intel i5, 22x Intel i7, 4x Tesla_K40), energy consumption between ~2 kW max and ~0.5kW on idle
Cooling System	Cooling Capacity: 10 kW (max: 11 kW), Max power consumption: 3.2 kW, EER = 3.2, CoP = 3.5, 1st rack is half air-cooled half liquid-cooled, 2nd rack is fully air-cooled
Photovoltaic system	80 PV panels with 20 kW peak power and energy storage of 75 kWh

The challenge in this case is to minimize the pilot DC operational costs with energy, considering variable characteristics such as the energy prices, renewable energy generation, efficient air conditioning management and computing load. By adjusting the power usage of the micro DC in certain periods, it is possible to emulate adaptation to variability of energy prices or renewable generation. The micro DC servers are managed by the SLURM queuing system for batch jobs and by OpenStack while the power management system allows for an accurate analysis of energy consumption of the racks and cooling system. In the next sub-sections, we show how scenarios 1 and 3 are addressed in Poznan Micro DC for providing electrical flexibility and exploiting the IT load migration cross DCs.

3.1.1. Scenario 1: Single DC Providing Electrical Energy Flexibility

In this test case, we evaluate the electrical energy flexibility of a Pilot DC offered by the use of photovoltaic renewable and associated energy storage. The flexibility services that may be provided are: (i) congestion management by decreasing the DC energy demand from the grid and (ii) trade energy in case of a surplus of RES. The test benefits from the power capping mechanism allowing to dynamically adjust the power level, and thus shifting the corresponding load execution to run when renewable energy is available.
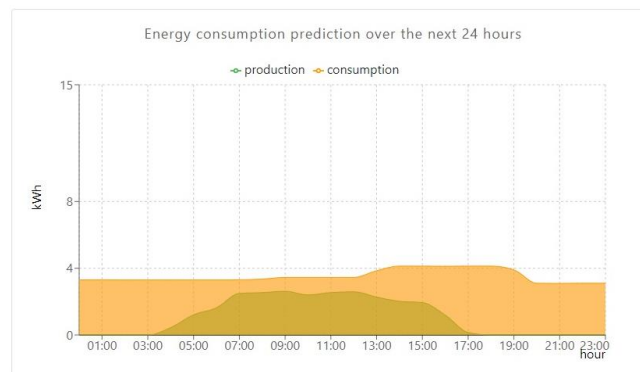
The DC Flexibility Manager System was deployed and configured into the pilot DC and used to manage the on-site electrical energy flexibility and PV utilization. It was integrated with the existing monitoring system which exposes PV production and energy data and stores them within the Data Storage. The micro DC energy consumption considers energy related to Real Workload, Delay Tolerant Workload and the cooling equipment. These data are used to forecast the energy demand for the next day, the potential latent energy flexibility and PV power production. The predicted data are later used by the Intra DC Energy Optimizer component to determine flexibility optimization plans shifting the electrical energy flexibility to maximize the locally generated renewable energy by leveraging on mechanisms such as workload management and battery storage usage. The provided action plan can be seen in DC Operator Console and is later translated into actions performed on the pilot DC.

The historical data for the pilot DC used by the energy prediction processes are shown in Figure 5 the DC energy consumption being split on the three main flexibility components (real-time workload, delay tolerant workload and cooling system).

Figure 5. Pilot DC monitored energy data displayed for 24 h.

The forecasts for the next operational day computed by the Electrical DR Prediction component are displayed in Figure 6.

Figure 6. Electricity prediction for the next day.

The Intra DC Energy Optimizer computes an energy flexibility management plan aiming to increase the DC demand during the day and decrease it during the night when the PV production is low. As shown in Figure 7a, the adapted DC energy demand (displayed as a dark green line) differs from the DC baseline energy demand (displayed as a dark line) because it will follow the pattern of the PV production, plotted as a light-green line. Thus, until the PV panels will start producing electricity, the plan aims at lowering the DC energy consumption by discharging batteries. They can support about 2 kWh of energy discharge each hour, until the capacity of 60 kWh of the battery depletes (light green columns in Figure 7b). Furthermore, the plan suggests shifting the delay tolerant workload from this period to the interval between hours 7 and 14 when the PV production is high. Each time the load shift action is encountered, the power is capped by the corresponding value and released later, according to the schedule. Moreover, the batteries will be recharged during the time period when the PV production is high (green columns in Figure 7b).

The optimization plan computed by the Intra DC Optimizer targets to leave the batteries at the end of the day in the same state as they were at beginning of the day. After the PV production stops, the DC energy consumption follows the normal pattern given by the baseline, since no actions are required.

The evaluated pilot DC is configured to be powered in equal parts by RES and the electricity grid. Moreover, the PV panels are used to charge the set of batteries. In the test, an assumption was considered to protect the batteries—not discharging them below 40%. Deeper (or worse—full) discharge leads to faster degradation of the battery and lower capacity. When there is a surplus of energy, solar cells charge the batteries; in the event of a shortage of energy, the batteries maintain the DC operation. Figure 8 depicts the charge level of the batteries during the test case execution.
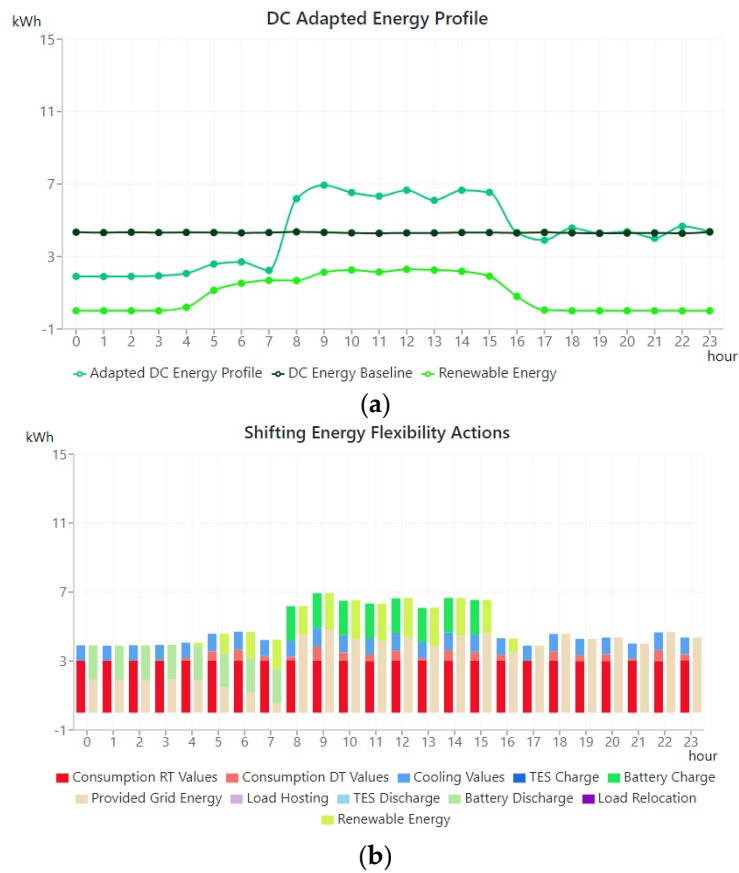
(a)

(b)

Figure 7. (**a**) DC adapted energy profile and (**b**) energy flexibility optimization plan.

Figure 8. Charging level of the batteries during the test case execution.

Charging and discharging the batteries takes place according to the schedule returned by the Intra DC Energy Optimizer. When it determines that the batteries are fully charged, it does not use them until start of the next cycle. However, some batteries power external devices that could not be turned off during the tests, hence the constant drop in the charge level after 6 p.m. in Figure 8. As mentioned, the testbed worked on real resources that were additionally used by external users. This may affect the appearance of additional tasks that did not alter the course of the test and confirms the possibility of its operation on an operational DC. Table 7 below shows relevant metrics calculated for this test case.

Table 7. Relevant metric values for this test case.

Metric	Formula	Description	Value
Renewable energy factor (REF)	$\frac{RE\ owned\ \&\ controlled\ by\ DC}{Total\ Facility\ Energy}$	% of renewable energy over total DC energy	0.428
Adaptability Power Curve at RES (APC$_{ren}$)	$1 - \frac{\sum_{i=1}^{n} \|K_{APCren} \cdot E_{Ren\ i} - E_{DCi}\|}{\sum_{i=1}^{n} E_{DCi}}$	Ability of a DC to adapt the energy demand (E_{DCi}) to the production curve of RES ($E_{Ren\ i}$).	0.28
Data Centre Adapt—DCA	$1 - \frac{\sum_{i=1}^{n}\|K_{DCA} \times E_{DCReal\ i} - E_{DCBaseline\ i}\|}{\sum_{i=1}^{n} E_{DCBaseline\ i}}$	Ability of a DC to change its energy consumption behaviour ($E_{DCReal\ i}$) considering the baseline consumption ($E_{DCBaseline\ i}$)	0.85
Grid Utilization Factor—GUF	$\frac{\sum_{i=1}^{n} f(n)}{N}$, $f(n) = \begin{cases} 1, & net\ exported < 0 \\ 0, & net\ exported \geq 0 \end{cases}$	% of time that locally RES generated energy cannot cover DC needs	0.66

3.1.2. Scenario 3: Workload Federated DCs

The goal of this test case was to exploit the possibility of IT load migration between DCs offered by the DCs Federation Manager system. We considered the pilot DC described the previous section as the destination DC. For the load migration, we utilized 16 servers Huawei XH620 blades, which reside within Huawei X6800 chassis. Each blade consists of 2x Xeon(R) E5-2640 v3 CPU (8 c/CPU) and 64 GB of RAM. The servers are connected to the Internet with symmetric 1 Gb/s connection. The DC deployed a DC Migration Controller Client around the OpenStack instance (Train distribution) hosted on a physical XH620 server with installed Ubuntu 18.04. OpenStack installation comes with Nova Compute and VPNaaS extension responsible for hosting the VM and providing VPN connections, respectively. The source DC provides with 32 servers equipped with AMD Ryzen 7 1700x and 32 GB of RAM with a similar software stack.
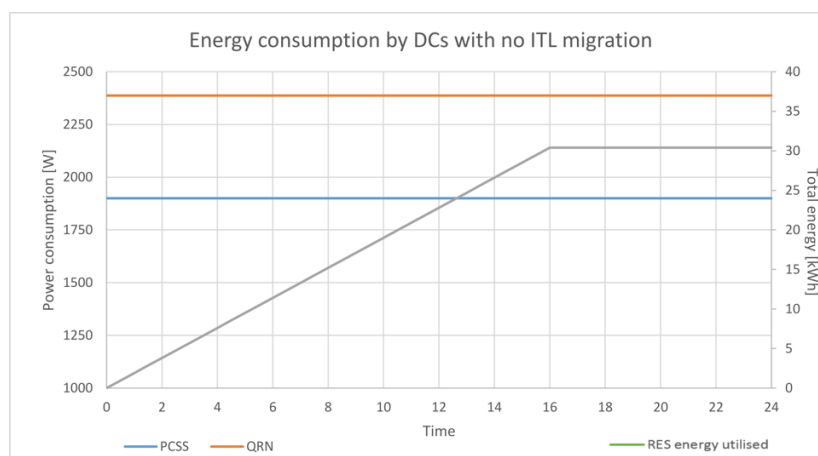
By the means of IT load migration, the amount of electrical and/or thermal flexibility of the destination DCs is increased and may provide to the local grid via DR programs or local flexibility markets. The migration test case was conducted between two servers of the source and destination DCs. The obtained results were used to extrapolate for the situation in which half of the workload is migrated to the destination DC (see Table 8).

Table 8. Measurements for test case and extrapolation.

Name of the Measurement	Value for a Single Server	Value for All Servers (Number in the Brackets)
Power usage of destination server in idle	92 W	1472 W (16)/736 W (8)
Power usage of destination server running single VM running Blender	125.3 W	(not used)
Power usage of destination server running 2 VMs running Blender	162 W	2592 W (16)/1296 W (8)
Power usage of destination server under the maximum load (or shifted PaaS load 2 servers to 1)	184 W	2944 W (16)/1472 W (8)
Power usage of destination server under typical (PaaS services) load	119 W	1900 W (16)
Power usage of source server running VM running Blender	74.6 W	2387 W (32)/1194 W (16)
Power usage of source server idle/switched off	20.5 W/xx	656 W/xx (32)/328 W/xx (16)
Migration duration of VMs between DCs (symmetric)	n/a	< = 22 min (16)/< = 26 min (32)

The destination DC servers' total idle power usage is estimated by multiplication of the value of energy drawn by a single server by the number of servers. A destination DC server can efficiently host two VMs running Blender. Such an assumption was driven by the fact that the destination DC node has 2x 8 core CPU with similar characteristics as the CPU of the source DC server and twice the amount of operational memory. The power drawn by the destination server hosting two VMs running Blender was computed as the sum of the power drawn by the idle destination server and twice the amount of power consumption that increased after a single VM migrated to the server. The latter component of the sum is enlarged by 5% in order to consider the worst-case scenario. The value of the power consumption was calculated by running a heavy load on one of the destination DC servers and measuring its power consumption. The sum for all servers was calculated as the value multiplied by 16, the number of servers. The assumption was made that after squeezing the PaaS services load running on 16 servers to eight servers, the power consumption of such servers would be much higher. In this experiment, a worst-case scenario was assumed, i.e., the maximum power drawn by such servers. A destination DC server under typical (PaaS services) load power usage was estimated as the average of the last 90 days.

In the first experiment, we considered that no workload was migrated between the source and destination DC (see Figure 9). This experiment was used as a baseline to the other two, since no load was shifted between the DCs. Not shifting the load made the power consumption in both DCs constant, and the RES energy utilization was the lowest among the experiments, i.e., 30.4 kWh.

Figure 9. Power consumption of servers where no workload is migrated from source to destination DC.

In the second experiment, half of the load was migrated in case of overproduction of RES at destination DC. The PaaS services running on destination servers were squeezed on half of the servers, inducing higher power consumption of the servers and resulting in lower QoS and reliability of the services. The other half of the servers were then released to receive the migrated workload. Figure 10 shows that workload from migrated resulted in power consumption decreased by half in the source DC. The power consumption at destination DC is higher due to the higher load of the PaaS services. This experiment allowed to draw approximately 36.5 kWh of energy from RES and is the most realistic due to proper PaaS services handling.
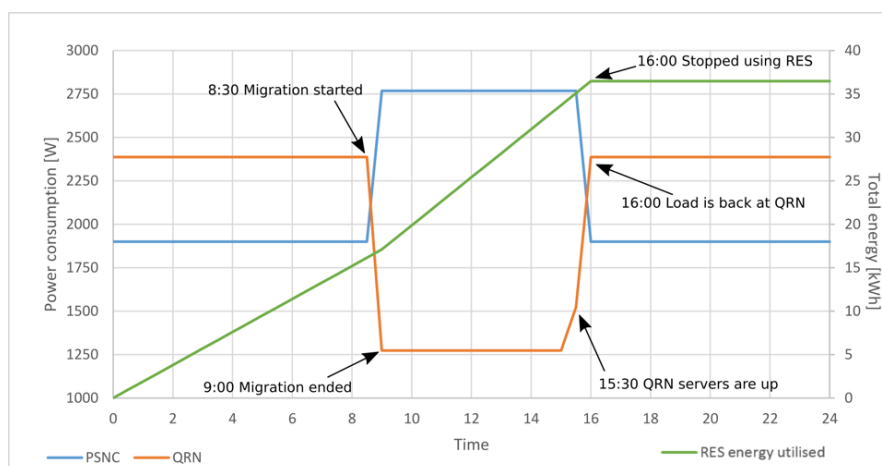
Figure 10. Power consumption of servers where half of the workload is migrated.

Table 9 presents the relevant KPIs calculated for this scenario. The REF and APCren metrics were only calculated for destination DC as source does not have access to renewable energy.

Table 9. Test case metrics values.

Experiment	KPIs	
	REF Dest. DC	APCren Dest. DC
No IT load is migrated from source DC to destination DC	0.67	0.40
Half of the IT load is migrated from source DC to destination DC	0.71	0.48

3.2. Pont Saint Martin DC

For these experiments, we leveraged on the pilot DC testbed provided by Engineering Informatica in Pont Saint Martin. The DC testbed cooling system is mainly composed of two Trane RTHA 380 refrigerating units (GFs): single-compressor, helical-rotary type water-cooled liquid chillers, accountable for server rooms' refrigeration and for the summer cooling of 6000 m^2 offices. Their characteristics are provided in Table 10.

Table 10. Characteristics of the refrigerant unit.

Refrigerating unit (GF)	Nominal Power: 100 kW, Tin_{C1} = 13 °C, $Tout_{C1}$ = 22 °C, Tin_{C2} = 10 °C, $Tout_{C2}$ = 7 °C, Coefficient of performance: 6.5

Moreover, each server room is equipped with two–three conditioning units (CDZs). The computer room air conditioner (CRAC) unit is composed of chilled water precision air conditioners systems, used to refrigerate the server rooms. They are located outside each server room: the air inside the server room flows through vents reaching the conditioning unit to be later reinjected into the bunker room through a ventilation shaft at the base of the rack's lines. Each server room includes a heterogeneous group of servers and devices, whose configurations differ: this means that each server-room will have a different number of hosted elements, power absorbed and therefore, set point temperature. Table 11 shows the characteristics of the sever rooms used for evaluation.

Table 11. Features of the server rooms used in evaluation.

Server Room	1	2
Ceiling height (m)	3.45	2.85
Area (m^2)	180.00	100.00
Gross volume (m^3)	621.00	285.00
Server room air set point temp. (°C)	23.50	24.50
Minimum Air Flow Temperature in Server Room	20.00	20.00
Maximum Air Flow Temperature in Server Room	26.00	26.00
Average IT power (kW)	96.00	44.60
Number of Racks	62.00	24.00
Number of servers only	330.00	154.00
Total number of device (server and storage) per Bunker	515.00	215.00
Number of cooling unit per bunker (GF id)	1	2
Total cooling power per unit (kW)	152.00	39.80
Nominal power required (kW)	20.00	n.a.
Nominal air flow pumped in the server room (m^3/h)	28,600.00	12,000.00

Finally, Table 12 presents the hardware characteristics of the other devices of the DC relevant for the experiments carried out.

Table 12. Heat pump and UPS system properties.

Device	Characteristics
Heat Pump	Coefficient of Performance (Cooling: 6.05, Heating: 4.43) Cooling mode: Power absorbed by the compressor = 131 KW Heating mode: Power absorbed by the compressor = 177.0 KW
UPS system	Power kVA: 800 kVA, Power kW: 640 kW, Battery type: EXIDE 14 OGi 800, Capacity: 1600 Ah, Number of cells: 204

In the next sub-sections, we show how scenarios 2 and 4 were addressed in this DC for re-using the generated heat and achieving more flexibility by efficiently combining power and heat.

3.2.1. Scenario 2: Single DC Providing Heat

The main goal of this experiment was to exploit the post cooling technique as a mechanism of thermal flexibility aiming to re-use the residual heat in nearby heat grids or buildings. We used the post cooling technique as a flexibility mechanism to increase the quality of the heat (i.e., temperature) recovered from the server rooms as much as possible without endangering the safe operation of the computing equipment.

The Flexibility Manager system was deployed in the pilot DC and used to compute the flexibility optimization actions. DCIM and Utility Networks Integration components were configured to allow on one side the gathering of real-time measurements from the pilot and the integration with local thermal grid for obtaining the heat demand and associated price to be used as flexibility optimization driver. Using the historical monitored data stored in the Data Storage, the Heat DR Prediction forecasts the day ahead energy demand and thermal flexibility of the DC. The Intra DC Energy Optimizer considers the heat prediction process results and heat demand to decide on an optimization action plan that contains post-cooling actions for specific intervals (i.e., change the configuration of bunker temperature alarm from 25 to 27 degree). The plans were calculated and used to shift thermal energy flexibility in response to specific heat reuse requests. Upon plan validation, the actions were executed including post-cooling of server rooms for more thermal flexibility by switching off refrigerator units GF1 and GF2 at specific time intervals.

The historical energy consumption data for the pilot DC is shown in Figure 11 (i.e., for the last 24 h), split on the three main flexibility components (i.e., real time workload, delay tolerant workload and cooling system). The DC energy consumption varies slightly around an average of 780 kWh.
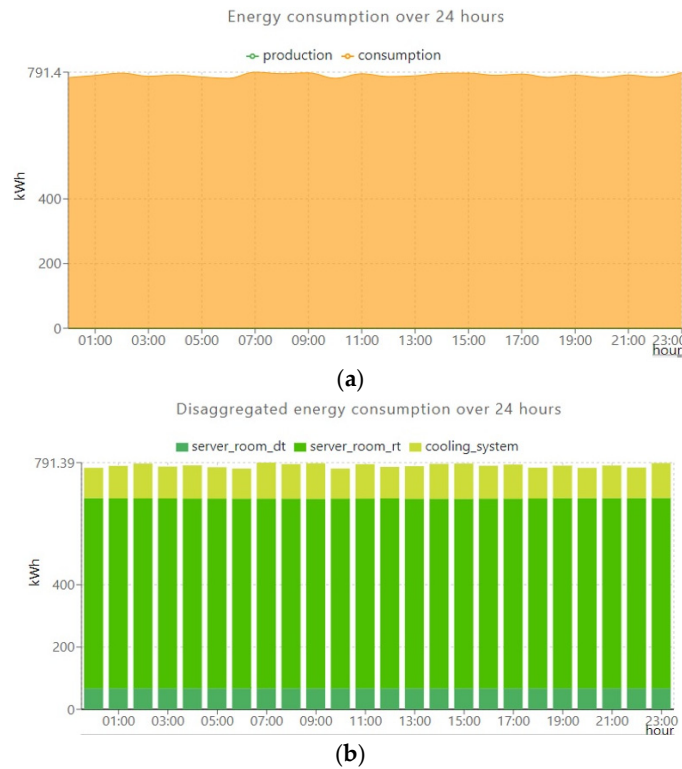
Energy consumption over 24 hours

(a)

Disaggregated energy consumption over 24 hours

(b)

Figure 11. Energy consumption data of the pilot DC (**a**-overall, **b**-disaggregated).

The Heat/Cold DR Prediction Module computes the estimation of the thermal generation for the next day using historical energy consumption data. The prediction over the next 24 h corresponds to the energy production of the last 24 h due to the low variability of the DC energy consumption.

The Intra DC Energy Optimizer uses the Marketplace Connector to get the heat demand and reference prices for the next day. As can be seen in Figure 12, the price for heat is high during hours 7–14; however, this does not match the thermal energy generation profiles as estimated by the prediction tool, which are almost constant during the day. Thus, the DC Optimizer computes an optimization plan to increase the DC thermal generation during the time interval when the heat price is high. This is done by leveraging on two flexibility mechanisms: shifting delay tolerant workload to increase the heat generation of the servers and post cooling the server room to extract more heat when the prices are high. The optimization plan is shown in Figure 13, in which the left chart shows the baseline heat generation (in dark green) and the adapted heat generation (in light green). Due to server room post cooling, the heat generation is increased during hours 8 and 9, enabling the DC to make extra profit by selling heat when the prices are high. The extra heat generated by the post cooling mechanism is shown as the dark blue columns in the right chart in Figure 13.
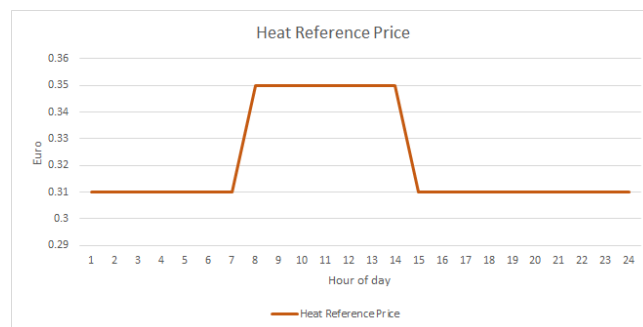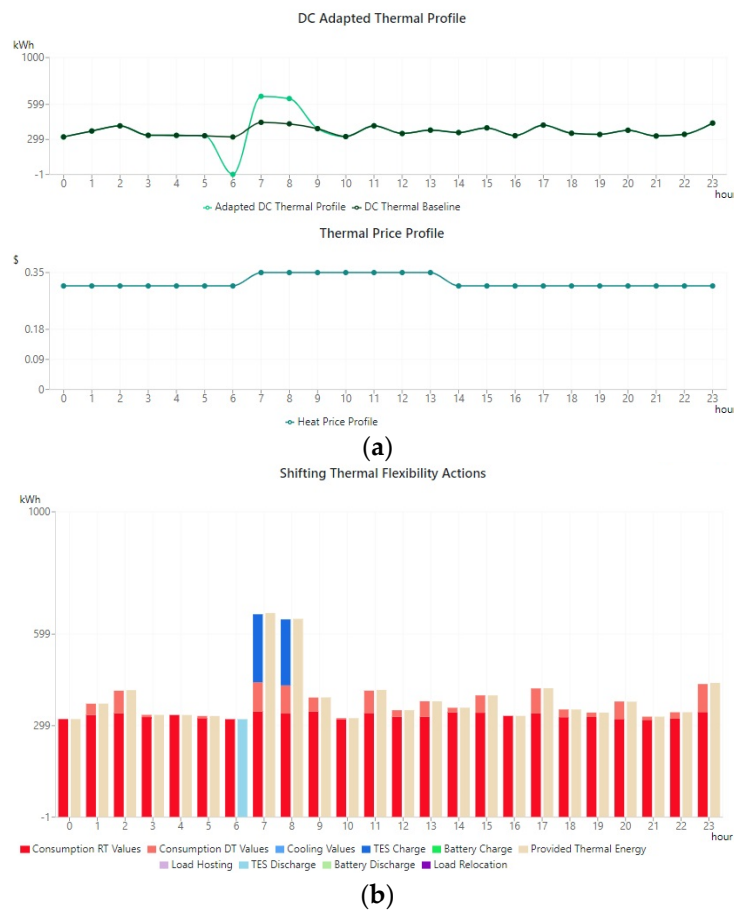

Figure 12. The heat reference prices considered in experiment.

Figure 13. Optimization plan from thermal perspective: (**a**) the DC thermal baseline and optimized profiles and thermal energy price, and (**b**) the DC energy flexibility shifting split on components.

Considering the execution of the day-ahead flexibility optimization plan, the metrics in Table 13 were calculated.
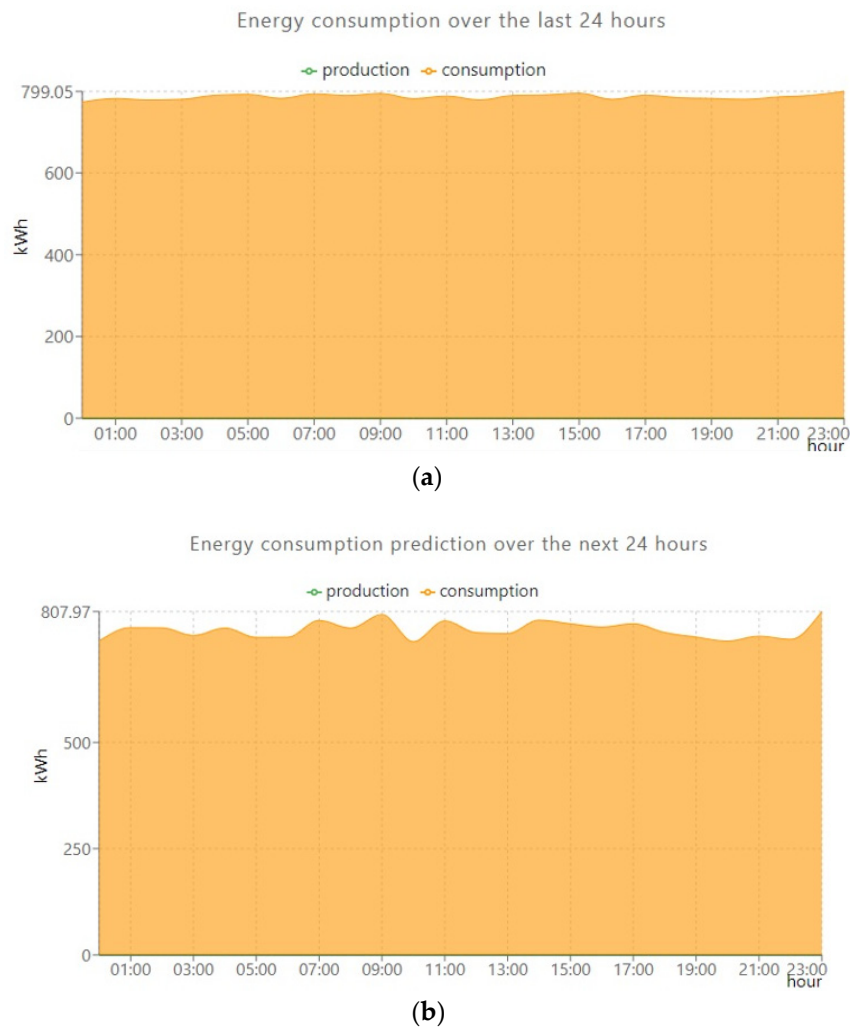
Table 13. Value of the metrics calculated for this experiment.

Metric	Formula	Description	Value
In House Reuse Factor (IRF)	$\frac{Energy\ Recovered\ within\ DC}{Total\ Facility\ Energy}$	% of the energy recovered within DC from the total facility energy	0.49
Energy Reuse Factor (ERF)	$\frac{Reuse\ X\ SourceFactor}{Total\ Facility\ Energy}$	% of energy that is exported for reuse outside of the data centre	0.48
Sustainable Heat Exploitation (SHE)	$\frac{electricity\ feeding\ the\ heat\ recovery\ system}{DC\ overall\ electricity\ consumption}$	% of electricity feeding the heat recovery system from the DC overall electricity consumption (before waste heat recovery)	0.14
Heat Usage Effectiveness (HUE)	$\frac{heat\ recovered}{SHE}$	heat recovered divide by SHE	2641.68

3.2.2. Scenario 4: Single DC Providing both Electrical Energy Flexibility and Heat

The goal of this experiment was to show how the pilot DC can operate at the crossroad of two energy networks (power and heat) and provide a combination of electrical and thermal flexibility. In this case, we leveraged on the possibility of controlling the cooling system and postponing the execution of delay tolerant workload as sources of flexibility. Using the monitored data stored in the Data Storage, the Heat DR Prediction and Electrical DR Prediction modules forecast the day ahead energy demand, electrical energy flexibility and thermal flexibility of the DC (see Figure 14). The Intra
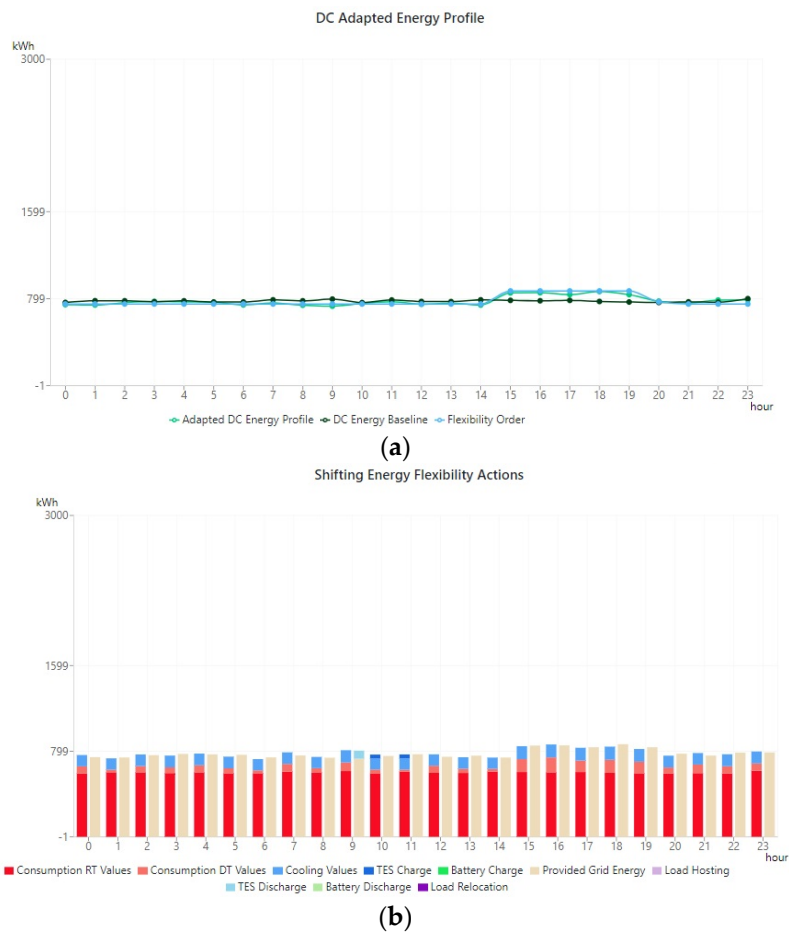
DC Energy Optimizer considers the electrical and heat flexibility prediction process results and both the electrical flexibility request and heat demand to decide on an optimization action plan. As a result, the optimization actions plans are calculated and used to shift both electrical and thermal energy flexibility to match specific flexibility requests.
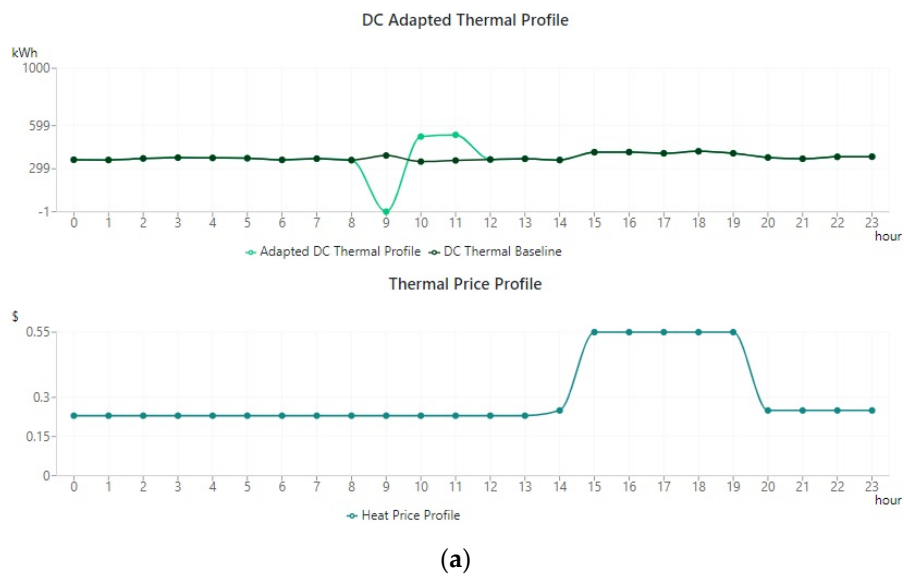

(**a**)

(**b**)

Figure 14. 24 hours energy (**a**) consumption and (**b**) prediction.

We leveraged on both shifting delay tolerant workload actions and post-cooling of the server room to allow the DC to adapt both its electrical energy consumption and the thermal energy generation to match the request. Figure 15 shows the DC optimization plan from the electrical energy flexibility perspective. The chart in the upper-left corner shows the DC baseline (in black), the electrical energy flexibility request having a peak between hours 15 and 20 (blue line) and the DC adapted profile that matches closely the Flexibility request (shown in green line). The DC energy consumption is shown in the upper right chart of its components.
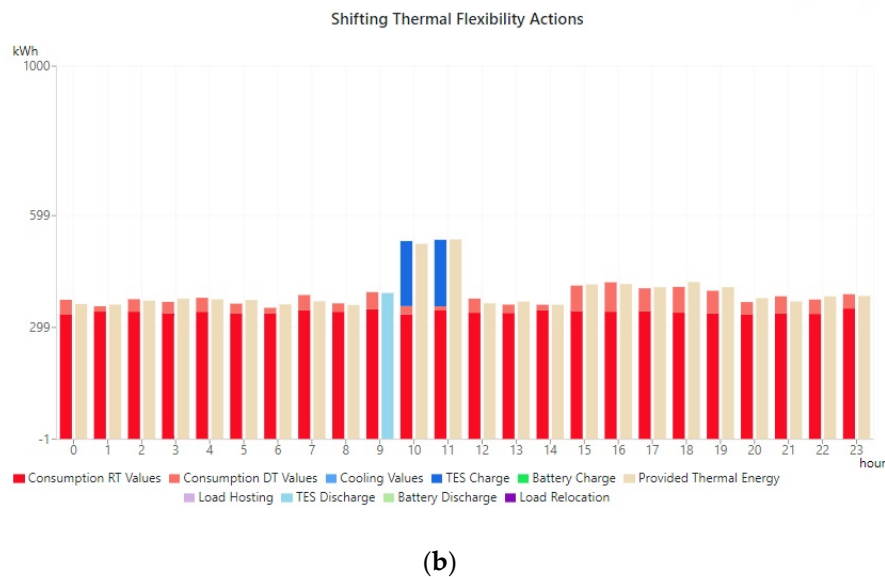
Figure 16 shows the same optimization plan from the thermal energy flexibility perspective. It shows the thermal flexibility of the DC and the generated heat that will be re-used in the local heat grid. The chart from the upper left corner shows the baseline heat generation (dark green) and the adapted heat generation (light green) and the effect of the post cooling action on the DC heat generation can be seen by the dark blue columns in the right chart.

(**a**)

(**b**)

Figure 15. Pilot DC electrical energy flexibility optimization plan: (**a**) the flexibility request and response, (**b**) the DC energy flexibility split on components.

(**a**)

Figure 16. *Cont.*

(**b**)

Figure 16. Pilot DC thermal flexibility optimization plan: (**a**) the DC thermal energy baseline, adapted profile and heat price, (**b**) the DC energy flexibility actions split on components.

For the day ahead flexibility optimization plans in this experiment, the metrics in Table 14 were calculated.

Table 14. Value of metrics calculated for this test case.

Metric	Description and Formula	Value
DCA	see Table 7	0.1
ERF	see Table 13	0.48
Primary Energy Savings (PES)	% of savings in terms of primary energy actually consumed by a DC and measured during period $\Delta t(PE_{Current,\Delta t})$ versus the baseline consumption ($PE_{Baseline\_adjusted_{\Delta t}}$); formula: $$1 - \frac{PE_{Current,\Delta t}}{PE_{Baseline\_adjusted_{\Delta t}}}$$	0.14
HUE	see Table 13	2650.44

4. Conclusions

In this paper, we address DCs' energy efficiency from the perspective of their optimal integration with utility networks such as electrical, heat and data networks. We describe innovative scenarios and ICT technology that allows them to shift electrical and thermal energy flexibility and exploit workload migration inside a federation to obtain primary energy savings and contribute to the grid sustainability.

The technology and several of the proposed scenarios were validated in the context of two pilot DCs: a micro DC in Poznan which has on site renewable and a DC in Point Saint Martin. The first experiment conducted on Poznan micro DC proved the possibility of using RES in DCs energy flexibility. The PV-system allows not only to reduce the grid power usage and reduce CO_2 emission, but also can be considered as a resource in energy flexibility services. By harnessing the energy storage, one may accumulate the energy and use it once more profitable. The profitability can be either driven by Congestion Management services or participation in DR.

The second experiment combined IT load migration with the availability of RES to increase the amount of energy flexibility and to find a trade-off between the flexibility level, QoS and the RES production level. As expected, migrating the load allows a reduction in CO_2 emission by opening possibilities for bolder power savings actions at DC originating the load. In this case, it is the effect of shutting down the unused nodes at one DC site and taking advantage of RES together with temporarily

postponing some of the running services on the other. However, the amount of workload relocated affects efficiency indicators in the opposite way. When the number of migrated VMs exceeds the computational capabilities of the destination DC, it may overload the system, thus, inducing less favorable exploitation of renewable energy. The key point is to find a balance between the power profile of DC and the acceptance level for the external load. One should easily note that the IT load migration can serve not only as a source of energy flexibility but can also lead to an increase in the overall energy efficiency of the federated system. In this way, it can match IT demands with time-varying onsite RES and implement "follow the energy approach".

In Point Saint Martin DC, the first performed experiment shows how the DC can adapt its thermal energy profile in order to match specific heat re-use requests. The obtained results prove that the heat recovery takes place even if the total facility energy is not being recovered for use within the DC boundaries. For efficient heat recovery, we considered the deployment of a heat pump as flexibility actions we have used the room post cooling by controlling the temperature set points and cooling system. The second experiment shows that both the electrical and thermal energy profiles of the DC can be adapted to match specific flexibility requests. The small value of primary energy savings obtained indicates that the plan focuses on adapting the DC energy demand and not reducing it while the reuse factor indicates that almost half of the energy consumed by the DC will be further reused in nearby destring heating systems.

Finally, to provide a guideline on the use of proposed ICT theology and applicability of the specific scenarios and flexibility actions, Table 15 shows that they fit to different DC types. We used a DCs' well known classification according to their specific operation, such as collocation, cloud and High-Performance Computing (HPC). For each DC type, we highlighted the possibility of applying a specific scenario considering their resources and flexibility sources.

Table 15. Suitability of specific scenarios and flexibility for data center types (+ suitable, - not suitable).

Scenarios		Data Center Type		
No.	Flexibility Source	Colocation	Cloud	HPC
	Delay-tolerant workload shifting	-	+	+
1	Use of cooling inertia	+	+	+
	Use of RES and energy storage	+	+	+
	Use of diesel generators	+	+	+
2	Local heat re-use	+	+	+
	Heat re-use at external entities	+	+	+
3	IT load migration	-	+	-
4	Dynamic usage of the cooling system and shifting of delay tolerant workload	-	+	-
5	IT load migration based on optimal DC power usage levels	-	+	-
6	IT load migration + heat re-use	-	+	-
7	IT load migration with energy flexibility and heat re-use	-	+	-

Generally, cloud DCs are suitable for most scenarios due their flexibility, moderate utilization, and good control. Of course, these are general guidelines and suitability depends strongly on a specific case, e.g., used software and technologies, data center policies, customers' requirements, etc. However, flexibility actions based on managing the workload are not well suited to collocation DCs that do not have direct control on servers and running workloads.

Conflicts of Interest: The authors declare no conflict of interest.

References

1. Rong, H.; Zhang, H.; Xiao, S.; Li, C.; Hu, C. Optimizing energy consumption for data centres. *Renew. Sust. Energ. Rev.* **2016**, *58*, 674–691. [CrossRef]

2. Cioara, T.; Anghel, I.; Salomie, I.; Antal, M.; Pop, C.; Bertoncini, M.; Arnone, D.; Pop, F. Exploiting data centres energy flexibility in smart cities: Business scenarios. *Inf. Sci.* **2019**, *476*, 392–412. [CrossRef]

3. Antal, M.; Pop, C.; Cioara, T.; Anghel, I.; Salomie, I.; Pop, F. A system of systems approach for data centers optimization and integration into smart energy grids. *Future Gener. Comput. Syst.* **2020**, *105*, 948–963. [CrossRef]

4. Wahlroos, M.; Pärssinen, M.; Manner, J.; Syri, S. Utilizing data center waste heat in district heating—Impacts on energy efficiency and prospects for low-temperature district heating networks. *Energy* **2017**, *140*, 1228–1238. [CrossRef]

5. Lis, A.; Sudolska, A.; Pietryka, I.; Kozakiewicz, A. Cloud Computing and Energy Efficiency: Mapping the Thematic Structure of Research. *Energies* **2020**, *13*, 4117. [CrossRef]

6. Antal, M.; Cioara, T.; Anghel, I.; Gorzenski, R.; Januszewski, R.; Oleksiak, A.; Piatek, W.; Pop, C.; Salomie, I.; Szeliga, W. Reuse of Data Center Waste Heat in Nearby Neighborhoods: A Neural Networks-Based Prediction Model. *Energies* **2019**, *12*, 814. [CrossRef]

7. Antal, M.; Cioara, T.; Anghel, I.; Pop, C.; Salomie, I. Transforming Data Centers in Active Thermal Energy Players in Nearby Neighborhoods. *Sustainability* **2018**, *10*, 939. [CrossRef]

8. Pärssinen, M.; Wahlroos, M.; Manner, J.; Syri, S. Waste heat from data centers: An investment analysis. *Sustain. Cities Soc.* **2019**, *44*, 428–444. [CrossRef]

9. Wahlroos, M.; Pärssinen, M.; Rinne, S.; Syri, S.; Manner, J. Future views on waste heat utilization—Case of data centers in Northern Europe. *Renew. Sust. Energ. Rev.* **2018**, *82*, 1749–1764. [CrossRef]

10. Huang, P.; Copertaro, B.; Zhang, X.; Shen, J.; Löfgren, I.; Rönnelid, M.; Fahlen, J.; Andersson, D.; Svanfeldt, M. 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. [CrossRef]

11. Data Centers Liquid Cooling. Novel Techniques for Optimal Thermal Flexibility Shifting and On-Demand Waste Heat Re-Use (CoolDC) Project. Available online: http://coned.utcluj.ro/cooldc/ (accessed on 15 October 2020).

12. Shao, W.; Chen, Q.; He, K.; Zhang, M. Operation Optimization of Liquid Cooling Systems in Data Centers by the Heat Current Method and Artificial Neural Network. *J. Therm. Sci.* **2020**, *29*, 1063–1075. [CrossRef]

13. Antal, M.; Pop, C.; Petrican, T.; Vesa, A.V.; Cioara, T.; Anghel, I.; Salomie, I.; Niewiadomska-Szynkiewicz, E. MoSiCS: Modeling, simulation and optimization of complex systems—A case study on energy efficient datacenters. *Simul. Model. Pract. Theory* **2019**, *93*, 21–41. [CrossRef]

14. Cioara, T.; Anghel, I.; Bertoncini, M.; Salomie, I.; Arnone, D.; Mammina, M.; Velivassaki, T.; Antal, M. Optimized Flexibility Management enacting Data Centres Participation in Smart Demand Response Programs. *Future Gener. Comput. Syst.* **2018**, *78*, 330–342. [CrossRef]

15. Vesa, A.V.; Cioara, T.; Anghel, I.; Antal, M.; Pop, C.; Iancu, B.; Salomie, I.; Dadarlat, V. Energy Flexibility Prediction for Data Center Engagement in Demand Response Programs. *Sustainability* **2020**, *12*, 1417. [CrossRef]

16. Koronen, C.; Åhman, M.; Nilsson, L.J. Data centres in future European energy systems—Energy efficiency, integration and policy. *Energy Effic.* **2020**, *13*, 129–144. [CrossRef]

17. Impram, S.; Varbak Nese, S.; Oral, B. Challenges of renewable energy penetration on power system flexibility: A survey. *Energy Strategy Rev.* **2020**, *31*, 100539. [CrossRef]

18. Kondziella, H.; Bruckner, T. Flexibility requirements of renewable energy based electricity systems—A review of research results and methodologies. *Renew. Sust. Energ. Rev.* **2016**, *53*, 10–22. [CrossRef]

19. Lombardi, P.A.; Moreddy, K.R.; Naumann, A.; Komarnicki, P.; Rodio, C.; Bruno, S. Data Centers as Active Multi-Energy Systems for Power Grid Decarbonization: A Technical and Economic Analysis. *Energies* **2019**, *12*, 4182. [CrossRef]

20. Basmadjian, R. Flexibility-Based Energy and Demand Management in Data Centers: A Case Study for Cloud Computing. *Energies* **2019**, *12*, 3301. [CrossRef]
21. Dabbagh, M.; Hamdaoui, B.; Rayes, A.; Guizani, M. Shaving Data Center Power Demand Peaks Through Energy Storage and Workload Shifting Control. *IEEE Trans. on Cloud Comput.* **2019**, *7*, 1095–1108. [CrossRef]
22. Anghel, I.; Cioara, T.; Salomie, I.; Copil, G.; Moldovan, D.; Pop, C. Dynamic frequency scaling algorithms for improving the CPU's energy efficiency. In Proceedings of the 2011 IEEE 7th International Conference on Intelligent Computer Communication and Processing, ICCP 2011, Cluj-Napoca, Romania, 25–27 August 2011; pp. 485–491. [CrossRef]
23. Cioara, T.; Salomie, I.; Anghel, I.; Chira, I.; Cocian, A.; Henis, E.; Kat, R. A Dynamic Power Management Controller for Optimizing Servers' Energy Consumption in Service Centers. In *Service-Oriented Computing*; Maximilien, E.M., Rossi, G., Yuan, S.T., Ludwig, H., Fantinato, M., Eds.; Springer: Berlin, Germany, 2010. [CrossRef]
24. Renugadevi, T.; Geetha, K.; Prabaharan, N.; Siano, P. Carbon-Efficient Virtual Machine Placement Based on Dynamic Voltage Frequency Scaling in Geo-Distributed Cloud Data Centers. *Appl. Sci.* **2020**, *10*, 2701. [CrossRef]
25. Pei, P.; Huo, Z.; Martínez, O.S.; Crespo, R.G. Minimal Green Energy Consumption and Workload Management for Data Centers on Smart City Platforms. *Sustainability* **2020**, *12*, 3140. [CrossRef]
26. Sajid, S.; Jawad, M.; Hamid, K.; Khan, M.U.S.; Ali, S.M.; Abbas, A.; Khan, S.U. Blockchain-based decentralized workload and energy management of geo-distributed data centers. *Sustain. Comput. Inform. Syst.* **2020**, 100461. [CrossRef]
27. Masdari, M.; Khoshnevis, A. A survey and classification of the workload forecasting methods in cloud computing. *Cluster Comput.* **2020**, *23*, 2399–2424. [CrossRef]
28. H2020 Catalyst Project. Available online: http://project-catalyst.eu/ (accessed on 15 October 2020).
29. Railis, K.; Voulkidis, A.; Velivassaki, T.H.N. Federated DC Migration Controller Software, GitLab. Available online: https://gitlab.com/project-catalyst/releases/federated-dc-migration-controller (accessed on 15 October 2020).