

We are IntechOpen, the world's leading publisher of Open Access books Built by scientists, for scientists

5,300

Open access books available

130,000

International authors and editors

155M

Downloads

Our authors are among the

154

Countries delivered to

TOP 1%

most cited scientists

12.2%

Contributors from top 500 universities



WEB OF SCIENCE™

Selection of our books indexed in the Book Citation Index
in Web of Science™ Core Collection (BKCI)

Interested in publishing with us?
Contact book.department@intechopen.com

Numbers displayed above are based on latest data collected.
For more information visit www.intechopen.com



Globally Optimised Energy-Efficient Data Centres

Dirk Pesch, Susan Rea, J. Ignacio Torrens,
Vojtech Zavrel, J.L.M. Hensen, Diarmuid Grimes,
Barry O'Sullivan, Thomas Scherer, Robert Birke,
Lydia Chen, Ton Engbersen, Lara Lopez, Enric Pages,
Deepak Mehta, Jacinta Townley and
Vassilios Tsachouridis

Additional information is available at the end of the chapter

<http://dx.doi.org/10.5772/65988>

Abstract

Data centres are part of today's critical information and communication infrastructure, and the majority of business transactions as well as much of our digital life now depend on them. At the same time, data centres are large primary energy consumers, with energy consumed by IT and server room air conditioning equipment and also by general building facilities. In many data centres, IT equipment energy and cooling energy requirements are not always coordinated, so energy consumption is not optimised. Most data centres lack an integrated energy management system that jointly optimises and controls all its energy consuming equipments in order to reduce energy consumption and increase the usage of local renewable energy sources. In this chapter, the authors discuss the challenges of coordinated energy management in data centres and present a novel scalable, integrated energy management system architecture for data centre wide optimisation. A prototype of the system has been implemented, including joint workload and thermal management algorithms. The control algorithms are evaluated in an accurate simulation-based model of a real data centre. Results show significant energy savings potential, in some cases up to 40%, by integrating workload and thermal management.

Keywords: energy efficient data centres, workload management, thermal management, integrated data centre energy management platform

1. Introduction

Data centres have become a critical part of modern information technology (IT) infrastructure with software as a service, mobile cloud applications, digital media streaming and the

expected growth in the Internet of Everything all relying on data centres. However, data centres are also significant primary energy users and now consume in the order to 3% of worldwide electricity and are responsible for 2% of global greenhouse gas emissions, the same as the airline industry [1]. With the increasing move towards cloud computing and storage as well as everything as a service type computing, data centre energy consumption is currently growing at a compound annual rate of over 10% and expect to reach approximately 8% of global energy consumption by 2020 [2, 3]. While the hyper-scale data centres of large cloud service providers are consuming in the 10 s of megawatts of power with corresponding annual electricity bills in the order of tens of millions of dollars, for example, Google with over 260 MW and \$67 M and Microsoft with over 150 MW and \$36 M in 2010 [4], those large cloud service providers are also investing heavily in energy efficiency and green data centres, for example, Google and Microsoft have invested over \$900 M in energy reduction measures since 2010. However, smaller operators, independent and co-location/multi-tenant data centres have not yet been able to deploy many of the energy efficiency technologies that are available. This is due to lack of integrated technology solutions and uncertainty about costs and the use of renewable energy solutions. In particular, the many server rooms and small data centres run by commercial businesses and universities are the dominant electricity users as shown in **Figure 1** [5].

On average, computing consumes 60% of total energy in data centres while cooling consumes 35% [6]. New server and cooling technologies have the potential to lead to a 40% reduction of energy consumption, but computation and cooling typically operate without joint coordination or optimisation. While server energy management can reduce energy use at CPU, rack and overall data centre level, dynamic computation scheduling is often neither efficient with many idle servers running rather than being shutdown [5] nor is it generally integrated with cooling. Data centre cooling typically operates at constant cold air temperature to protect the hottest server racks, while local fans distribute the air across racks. However, these local




Figure 1. Estimated US data centre electricity consumption by market segment (2011) [5].

server controls are typically not integrated with room cooling systems, which means that it is not possible to optimise chillers, air fans and server fans as a single, whole system.

In order to reduce the CO₂ footprint of data centres, large organisations such as Google and Facebook are investing in renewable energy sources (RES), such as solar photovoltaics (PV) or wind power, often co-located with their hyper-scale data centres [7, 8]. However, for the many smaller data centres and server rooms, the use or integration of renewable energy sources has received limited interest. The reason for this is that these data centres are typically embedded in buildings that also hold other functions, for example, office and meeting spaces, laboratories and lecture rooms in the case of universities. A major issue in this is also the lack of interoperability of generation, storage and heat recovery and current installation and maintenance costs versus payback [9]. By and large, data centre operators, who want to be green and use renewable energy, buy electricity that has been given a green label by their respective supplier without often being able to fully verify this. The intermittency of renewable energy generation is also a critical factor in an environment with very strict service level agreements and essentially 100% uptime requirements. The adoption of new technologies related to computing, cooling, generation, energy storage and waste heat recovery individually requires sophisticated controls, but no single manufacturer provides a complete system, so integration between control systems does not exist.

However, research has been under way in a cluster of projects funded by the European Commission's Framework Programme for Research and Innovation. The cluster includes projects such as DC4Cities, GENiC, CoolEmAll, RenewIT, Eureca, GEYSER, GreenDataNet, Dolfin and All4Green, which are all focused on a range of aspects to increase data centre energy efficiency but also to integrate data centre energy use and recovery into a future smart grid and smart city environment. One of those projects, GENiC (<http://www.projectgenic.eu>), in particular, aims at developing integrated cooling and computing control strategies in conjunction with innovative power management concepts that incorporate renewable electrical power supply and storage, and waste heat management. The project's aim is to address the issue mentioned above by developing an integrated, flexible, component-based management and control platform for data centre wide optimisation of energy consumption, reduction of carbon emissions and increased local renewable energy supply usage through integrating monitoring and control of computation, data storage, cooling, on-site power generation and waste heat recovery.

A key element in not only achieving a reduction in energy consumption but also a reduction in carbon emissions is energy supply by renewable energy generation and, where possible, energy storage equipment. Such an approach needs to be operated as a complete system to achieve an optimal energy and emissions outcome. This vision of integrated, holistic energy management is centred on the development of a hierarchical control system to operate all of the primary data centre components in an optimal and coordinated manner.

2. Challenges for integrated data centre energy management

While data centres have become a critical IT infrastructure and also a significant consumer of energy and contributor to CO₂ emissions, opportunities exist to enhance the energy and power

management of data centres in conjunction with renewable energy generation and integration with their surrounding infrastructure. Work has been done on studying the topic of powering of data centres by renewable energy [10], but this has not been fully integrated into a complete energy management system considering coordinated workload management, cooling, powering and heat recovery management. While much work has focused on integrated energy management for data centres [11, 12], there is still a lack of an overall consideration of energy usage and powering with the recovery of waste heat as part of an overall thermal management approach. In order to bring the elements of workload management, cooling, powering and heat recovery together in such a way that it will be possible to achieve a high level of renewable energy powering of data centres, a comprehensive integrated energy management system is needed. The challenges that such a system needs to address are as follows:

- **Comprehensive, scalable integration of workload management with cooling approaches:** in most data centres, workload is allocated to servers without consideration of the thermal impact that this has on the data centre space. In many cases, idle servers are not even shut-down and continue to consume energy without any productive IT load processing. An integration of IT workload management with cooling through thermally aware workload consolidation is required.
- **Effective power management with a high level of renewable energy supply integration while meeting service level agreements:** in order to facilitate the uptake in renewable energy supply systems, in particular at a local level, intelligent power management approaches are needed to balance the intermittently available renewable energy sources, for example, solar, wind, with grid supplied electricity while managing service level agreements. Power management needs to also take energy price fluctuations and demand response requirements into consideration to maximise the cost-effectiveness of renewable power solutions in order to create incentives for investment in such solutions.
- **Strategies for waste heat recovery in conjunction with the heating needs of surrounding areas:** opportunities exist for small- to medium-sized data centres to reuse the heat generated by IT workload in order to heat adjacent spaces rather than dump the heat into the air through heat exchangers or dry coolers. Heat recovery solutions can heat spaces or water either within the same building or for larger data centres spaces in adjacent buildings or feed heat into local district heating systems. In this way, heat recovery can reduce the energy demands of adjacent facilities and achieve an overall reduction of energy consumption within the area of the data centre.
- **Design and decision support tools assisting data centre operators with data centre energy management:** for many data centres, in particular for small- to medium-sized data centres embedded into larger organisations, the IT manager and the facilities manager are different roles and as such do not have complete understanding of the complete energy management needs and opportunities. As such, suitable tools are required to assist operators with decision-making in terms of what energy management approaches, power solutions or heat recovery techniques might be most suitable for their situation.
- **Effective monitoring and fault management:** maintain service level agreements and uptime is of paramount importance to data centre operators, above and beyond of everything else.

In order to achieve this while making sure energy consumption costs do not exceed certain levels, effective monitoring and fault management tools are important and can assist operators with their work.

3. An architecture for globally optimised energy management in data centre

To address the challenges outlined above, the EC-funded GENiC project has developed a high-level architecture for an integrated design, management and control platform, targeting data centre wide optimisation of energy consumption by encapsulating monitoring and control of IT workload, data centre cooling, local power generation, energy storage and waste heat recovery. The developed management platform includes control and optimisation, decision support, and fault detection functions and defines interfaces and common data formats to enable a component-based design. The GENiC architecture can act as a template for a wide range of implementations of data centre energy management systems suited to a particular data centre configuration. In the following, a functional specification of the GENiC architecture is presented and an overview of the integration framework is provided. The applicability of the proposed functional architecture is illustrated by a number of use cases. More detail can be found in [13].

3.1. Functional architecture

The GENiC architecture integrates workload management, thermal management and power management by using a hierarchical control concept that enables the coordination of the management sub-systems in an optimal manner with respect to the cost of energy consumption, environmental impact and cost policies. **Figure 2** provides an overview of the developed GENiC system architecture, which consists of six functional groups, the GENiC component groups (GCGs):

- The **Workload Management GCG** is responsible for monitoring, analyzing, predicting, allocating and actuating IT workload within the data centre.
- The **Thermal Management GCG** is responsible for monitoring the thermal environment and cooling systems in the data centre, predicting temperature profiles and cooling demand, and optimally coordinating and actuating the cooling systems.
- The **Power & RES Management GCG** is responsible for monitoring and predicting power supply and demand, and for actuating the on-site power supply of the data centre.
- The **Supervision GCG** includes the supervisory intelligence which provides policies to the power, thermal and workload GCGs for supplying electrical power to meet the IT and cooling power demands of a DC based on monitoring data, predicted systems states and actuation feedback.
- The **Support Tools GCG** includes a number of tools that provide decision support for data centre planners, system integrators and data centre operators.




Figure 2. Overview of the GENiC architecture (from [13]).

- The **Integration Framework GCG** provides the communication infrastructure and data formats that are used for interactions between all components of the GENiC system.

Each GCG is composed of a number of functional components, the GENiC components (GCs) (see **Figure 2**). The core function of the GENiC system for continuous data centre energy optimisation can be divided into four basic steps:

1. **Monitoring** components within the management GCGs collect data about IT workload, thermal environment, cooling systems, power demand and on-site power supply.
2. **Prediction** components within the management GCGs update their internal models and estimate future system states based on the collected monitoring data.
3. **Optimisation** components determine optimal policies based on the collected monitoring data and calculated prediction data. These policies are provided to the management GCGs.
4. **Actuation** components within the individual management GCGs implement the policies provided by the optimisation components in the data centre and at the renewable energy sources facilities.

These elements are complemented by components for external data acquisition and fault detection and diagnostics. The basic information flow for coordinating workload, thermal and power management is illustrated in **Figure 3**. In the following, the GENiC component groups are described in more detail.

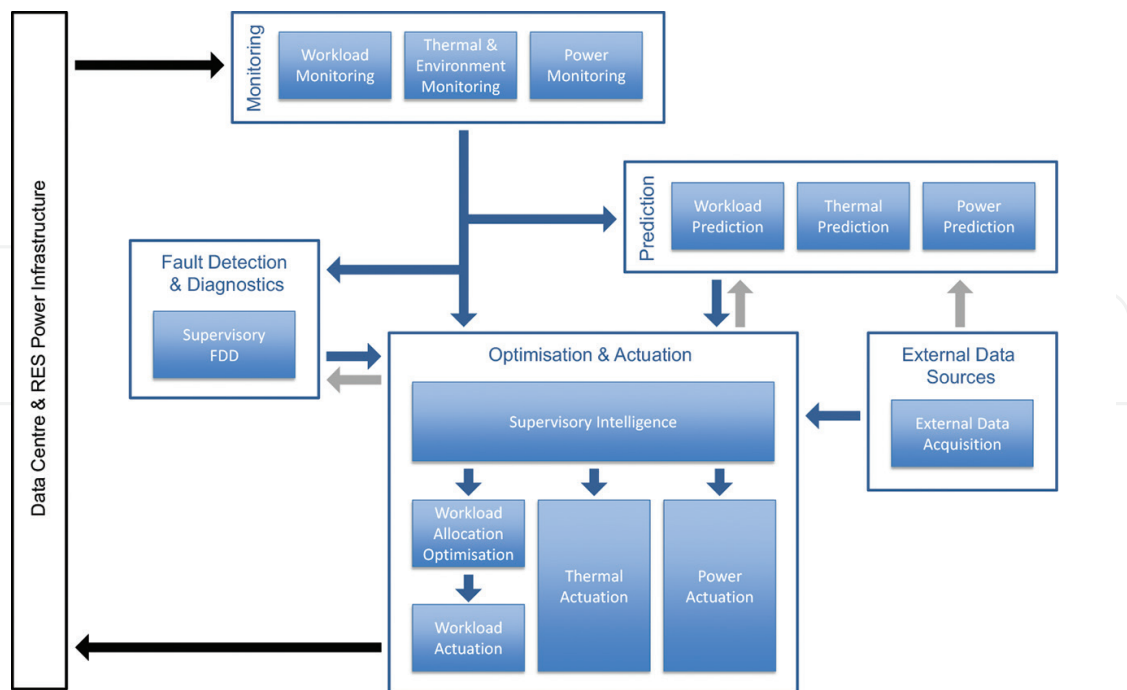


Figure 3. Information flow (simplified) for coordinating workload, thermal and power management [14].

Workload Management GCG: The primary objective of this GCG is to allocate virtual machines (VMs) to physical machines (PMs) such that service level objectives (SLOs) are satisfied with low operational cost. Monitoring data from the IT resources deployed within the data centre are collected by the Workload Monitoring GC. The Workload Prediction GC uses this information to provide short- and long-term predictions on resource utilization. The allocation and migration of VMs to PMs are determined by the Workload Allocation Optimisation GC, which solves a constrained optimisation problem, taking the predicted workload as well as constraints provided by the Supervisory Intelligence GC, Thermal Prediction and Performance Optimisation GC into consideration. The Performance Optimisation GC defines location constraints for individual VMs and modifies the individual VMs' priorities to fulfil application specific SLOs. The VM allocation plan is finally applied by the Workload Actuation GC, which provides an interface to the data centre-specific virtualization platform.

Thermal Management GCG: The Thermal & Environment Monitoring GC integrates monitoring of cooling systems and a sensor network infrastructure for collecting temperature and other environmental data in the data centre space. The collected data are used by the Thermal Prediction GC to provide short-term and long-term predictions to support supervisory control decisions, thermal actuation and workload allocation. Long-term predictions are used for making decisions at the supervisory level. Short-term thermal predictions are required by the Thermal Actuation GC along with real-time sensor measurements to determine optimal set

points for the cooling system in order to achieve the targets set by the Supervisory Intelligence GC. These short-term thermal predictions are also necessary input to the Workload Allocation Optimisation GC, as they include temperature models for the thermal contribution of IT server workload to the server inlets and the Supervisory Intelligence GC. Furthermore, short-term predictions, combined with equipment fault information from the Thermal Fault Detection & Diagnostics (FDD) GC, are used for fault detection and diagnostics at the supervisory level.

Power & RES Management GCG: The Power Monitoring GC provides power monitoring information of the DC (power consumed per server, per rack level and total DC power demand), as well as integrates monitoring of the RES infrastructure for local energy generation and storage with data centre power consumption requirements. These data are used by the Power Prediction GC to provide IT Power prediction as well as long-term predictions to support supervisory control decisions and power actuation. The Power Actuation GC determines operation set points for the power systems based on operation policies provided by the Supervisory Intelligence GC and adjusting them depending on measured data and operational conditions.

Supervision GCG: The Supervisory Intelligence GC is responsible for the overall coordination of workload, thermal, power management and heat recovery. It considers power demand and supply, grid energy price, energy storage availability and determines how much power should be supplied from the electricity grid, RES and energy storage to achieve a particular objective on power usage. To this end, it provides policies for the components in the Workload Management, Thermal Management and Power & RES Management GCGs based on information from monitoring and prediction components. The Supervisory Intelligence GC provides these high-level policies for the purpose of guiding the individual management functions towards the Supervisory Intelligence objective strategy that has been chosen as the driver for current data centre operations. Key objective choices might be minimization of financial cost, minimization of carbon emissions or maximization of RES usage. To detect and diagnose system anomalies, the Supervisory FDD GC compares predicted values with measurement data and collects and evaluates fault information. In appropriate situations, the Supervisory FDD GC informs the Supervisory Intelligence GC when a deviation becomes substantial enough to negatively impact system operation so that mitigation action can be taken by the platform until the fault has been corrected. The Human-Machine Interface GC provides a framework for user interfaces that allow data centre operators to monitor and evaluate aggregated data provided by the individual GCs.

Integration Framework GCG: The Communication Middleware GC provides the communication infrastructure used within the GENiC platform. The Data Centre Configuration GC uses a centralized data repository to store all information related to the data centre configuration, including information on data centre layout, cooling equipment, monitoring infrastructure, IT equipment and virtual machines running in the data centre. Finally, the External Data Acquisition GC provides access to data not collected by existing components of the GENiC platform, including weather data, grid energy prices and grid energy CO₂ indicators.

The GENiC platform integrates distributed software components, which are developed and maintained by individual consortium partners. A software component can implement a single

GC, multiple GCs or just part of a GC to provide the required functionality to the platform. For example, a topic-based publish-subscribe messaging architecture is a suitable mechanism to ensure a robust data exchange between individual software components. With this approach, the components do not need to be connected directly to each other, but components can publish messages to a central message broker using pre-defined topics and subscribe to the broker to topics from other components that are of interest to them. The broker forwards all incoming messages to the appropriate subscribers. The GENiC architecture defines a consistent interface specification using a common data format for all GENiC components. All interfaces are defined by hierarchically structured topics. Each of these topics has a defined message payload structure that uses the GENiC common data exchange format which is specified based on JSON [15]. This approach creates a very flexible data centre management platform that can be configured to suit individual, local data centre configurations.

Support Tools GCG: The GENiC platform includes a number of tools to assist data centre planners, system integrators and data centre operators:

- The Workload Profiler GC consists of a set of tools to capture application profiles that can be used by data centre operators to improve application performance.
- The Decision Support for RES Integration GC is a tool for data centre planners to determine the most cost-efficient renewable energy systems to install at a data centre facility.
- The Wireless Sensor Network (WSN) Design Tool GC is a tool to capture system and application level requirements for data centre wireless monitoring infrastructure deployments.
- The Workload Generator GC provides recorded and synthetic VM resource utilization traces for the simulation-based assessment of a GENiC-based system and its implemented algorithms and policies.
- The Simulator GC supports the testing of individual and groups of GCs as well as the (virtual) commissioning of a GENiC platform before its deployment in an actual data centre.
- The Multi Data Centre (DC) Optimisation GC is a tool that exploits the differences in time-zones, energy tariff plans, outside temperatures, performances of geographically distributed data centres to allocate workload amongst them in order to minimise global energy cost and related metrics.

3.2. Energy management use case

The GENiC project's focus to optimally operate data centres with respect to energy is achieved through the integration of workload management, thermal management and power management (including powering through renewable energy sources) via a hierarchical supervisory control concept. Key optimisation criteria in consideration by data centre operators are (i) meeting agreed service level agreements (SLAs), (ii) minimisation of total energy costs, and (iii) with the availability of renewable energy sources also, the maximisation of RES power use and minimisation of carbon emissions. To account for fluctuations in the IT workload demand and the availability of renewable energy supply (which includes local on-site energy production and grid power), the set points of the management sub-systems have to be

adapted over time. The Supervisory Intelligence (SI) GC coordinates the individual management sub-systems, including renewable energy supply, by providing optimal policies with respect to the selected optimisation criterion. The use case scenario is illustrated in **Figure 4**. The basic operational flow is as follows [14]:

Step 1—The monitoring GCs, Workload Monitoring, Thermal & Environment Monitoring, and Power Monitoring, collect data from VMs, PMs, air conditioning equipment, sensor networks, power meters and on-site energy supply systems. The relevant information is forwarded to the individual prediction and actuation GCs and SI.

Step 2—Based on recent and historical monitoring data, the prediction GCs, Workload Prediction, Thermal Prediction, and Power Prediction, predict server power demand, thermal profile and cooling demand, RES production capacity and energy demand. The relevant information is forwarded to the individual actuation GCs and SI.

Step 3—Additional data, that is, weather data and grid energy prices, are obtained from external data sources and forwarded to SI by the External Data Acquisition GC.

Step 4—SI provides a set of policies to the actuation GCs, Workload Allocation Optimisation, Thermal Actuation and Power Actuation that are based on inputs from the monitoring and prediction components and further interactions with the Power Prediction GC. These interactions validate the consequences of particular power profiles that SI considers as part of the policy definition. The Workload Allocation Optimisation GC solves a constrained optimisation

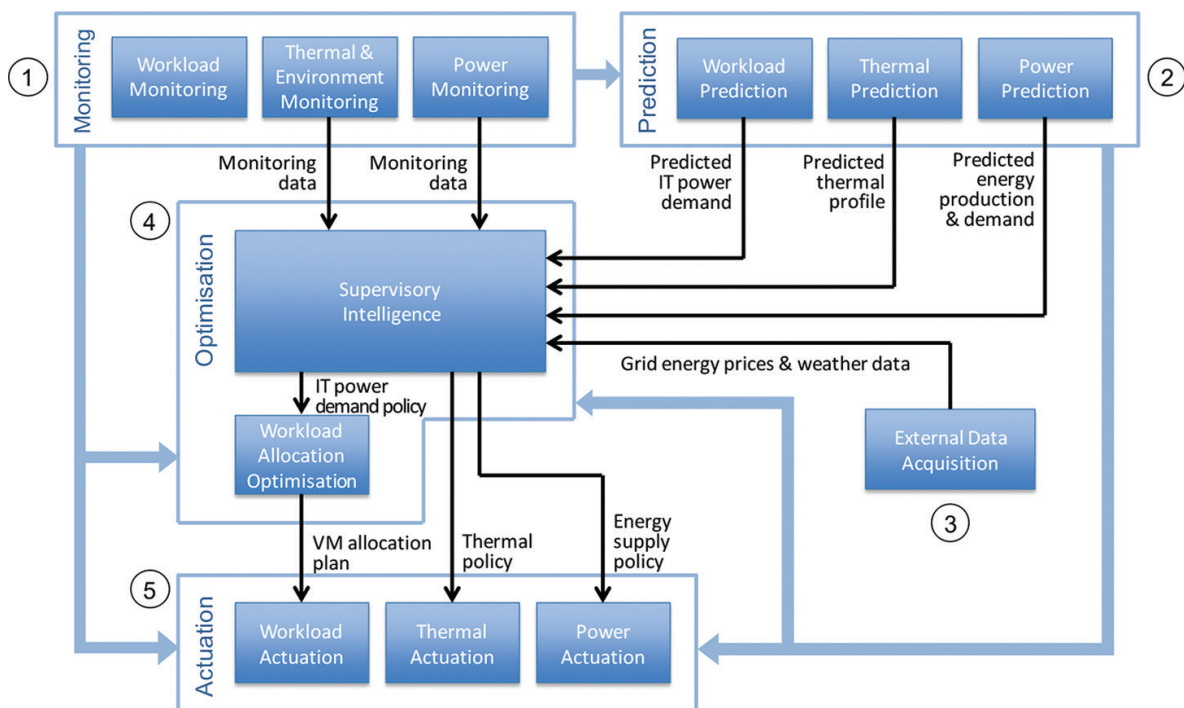


Figure 4. Energy management use case [14].

problem to determine an optimal VM allocation plan minimizing server energy consumption, taking the upper-bound IT power budget recommended by SI and additional inputs from other GCs (thermal and colocation and anti-colocation constraints) into consideration. The Thermal Actuation GC takes the minimum and maximum allowable data centre temperatures determined and then provided to it by SI and optimally calculates cooling equipment set points that ensure the room's thermal profile is properly regulated with minimal cooling equipment electrical power consumption. The Power Actuation GC implements the distribution plan for drawing electricity from grid, controllable and uncontrollable RES, and the schedule for charging and discharging the energy storage device.

Step 5—Based on the inputs from SI and the Workload Allocation Optimisation GC, as well as monitoring and prediction components, the actuation GCs, Workload Actuation, Thermal Actuation, and Power Actuation, decide and apply the actual control actions. For example, the Workload Actuation GC executes the VM allocation plan and switches PMs on/off, based on the actuation requests. Faults are reported back to the optimisation GCs to be considered in the next iteration of the optimisation process.

4. Prototype implementation

Figure 5 illustrates a prototype implementation of the GENiC architecture. The GENiC distributed architecture approach with clearly defined interfaces simplifies integration of a diverse set of software components and allows flexible configuration of the platform. Due to the diverse set of technologies in use in data centres, for example, IT systems, cooling systems, power systems and RES facilities, there is typically no individual manufacturer who supplies all the systems that a data centre requires. Therefore, a data centre management system architecture needs to allow for the integration of individual components supplied by multiple manufacturers and service providers. The architecture detailed in Section 3 is scalable and flexible at the same time and is based on micro-service architecture principles that offer the following benefits:

- **Separation of concerns**—each service implements a single operational functionality. The architecture becomes more flexible and scalable at the same time.
- **Distributed security compliance**—each service can have different security policies, allowing each service provider to maintain local security policies.
- **Freedom of service implementation**—each service provider can choose any development language without compromising the integrity of the overall platform. The only requirement is that the service needs to be able to communicate with the messaging broker.
- **Service scalability**—new instances of services can be spawned when more processing power is required.
- **Simplified API**—all modules use a common API to exchange data and trigger events used by other services.

- **Simplified testing and integration**—testing and integration are easier as testing focuses on black box testing with implementation details hidden behind APIs. Service integration hides APIs and dependencies.

A central element of the implementation of the prototype is the use of the RabbitMQ messaging system [16] for the exchange broker. RabbitMQ provides a range of client implementations in a wide range of programming languages, which allows manufacturers to suit their individual technology set-ups. A Generic Client architecture has been developed to allow each component provider expose their components in a distributed manner in the architecture. The individual GENiC components are implemented as services that communicate via the message broker. The client architecture also offers an easy way to integrate 3rd party (closed source) services with a minimal effort. Each of the components implemented in the GENiC prototype are shown in **Figure 5**, colour coded based on the component group they belong to. Short-term monitored data are stored in a database backend in the GENiC prototype implementation. CouchDB as a NoSQL solution is used, but many other data base solutions are possible depending on the specific needs and data volumes of a particular configuration. Due to the large quantity of stored data, only short-term data are available on the broker.




Figure 5. GENiC architecture implementation prototype.

5. Assessment of energy efficiency

In order to assess the effectiveness of data centre management systems in terms of the energy efficiency, power management, managing increased penetration of renewable energy sources, heat reuse and data centre flexibility, the need to select appropriate metrics is of paramount importance. The aforementioned cluster of European research projects on data centre energy efficiency has taken five common data centre metrics and defined 21 new metrics, along with measurement methodologies, to adequately capture the energy efficiency, flexibility and sustainability of modern data centres [17]. This approach supports the development of a common framework for monitoring and assessing the flexibility and sustainability of data centres. The metrics of specific interest for the evaluation of an integrated energy management platform, which integrates thermal and workload management with renewable energy/power supply and heat recovery, are listed in **Table 1**.

The GENiC project considers two types of evaluation: one is based on simulation-based assessment (SBA), which uses the Simulators GENiC component (see **Figure 2**), provided by the tools that have been developed in the project. The Simulators component provides a virtual data centre based on TRNSYS model implementation and simulation and additional interfacing and timing functions [18]. The SBA uses the full energy management platform in the same manner as it is used in a real physical data centre. SBA has the advantage that a specific architecture configuration can be tuned to a particular data centre set-up before deployment in the real environment. This allows for a priori energy efficiency assessment, which not only enables data centre operators to understand what energy savings can be expected from a deployment of an integrated data centre energy and power management platform, but also prepares the platform to run optimally once deployed without affecting the real environment during an in situ tuning process.

Metric	Goal
PUE—Power Usage Effectiveness	Energy/Power Consumption
CER—Cooling Effectiveness Rate	
CUE—Carbon Usage Effectiveness	
Energy Effectiveness of Cooling Mode in a Season	
ERE—Energy Reuse Effectiveness	Energy Recovered/Heat Recovered
APCren—Adaptation of Data Centre to Available Renewable Energy	Data Centre Flexibility—Energy Shifting
DCA—Change in Data Centre Energy Profile from Baseline	
RenPercent—Share of Renewables in Data Centre Electricity Consumption	Renewables Integration
Renewable Energy Factor	
CO ₂ Savings Change in Data Centre CO ₂ Emissions From Baseline	Primary Energy Savings and CO ₂ avoided emissions

Table 1. GENiC evaluation metrics.

The second evaluation is based on the deployment of the prototype in a real data centre. The project chose a small but typical data centre at Cork Institute of Technology. The data centre was adapted to the needs of the project to enable extensive control of the thermal management side, including heat recovery and both virtualisation of the computing infrastructure and normal operation. Experimental renewable energy facilities are linked in a virtual manner to the data centre as the renewable energy micro-grids are located on two premises of project partner Acciona in Spain. The demonstration of use of renewable energy is possible by recording the amount of energy that can be generated by typical micro-grids over time and accounting the amount of electricity flowing into the data centre as either non-renewable or renewable.

5.1. Simulation model—virtual C130 data centre

In order to evaluate the performance of the GENiC platform and to allow pre-deployment assessment and tuning, the project has developed a Simulators GC, which is part of the Support Tools GCG. The simulator component includes energy models that emulate the performance of a data centre and its systems, supporting the development and testing of GENiC components as well as the commissioning of the overall GENiC platform, prior to its physical deployment to the real data centre [19]. The Simulators GC consists of energy models shown in **Figure 6**. These are on the demand side, for example, data centre environment (building energy model and building airflow model), IT devices model, and heating, ventilation and air conditioning (HVAC) systems model, and the supply side, for example, power supply model.




Figure 6. Types of energy models in the Simulator GC.

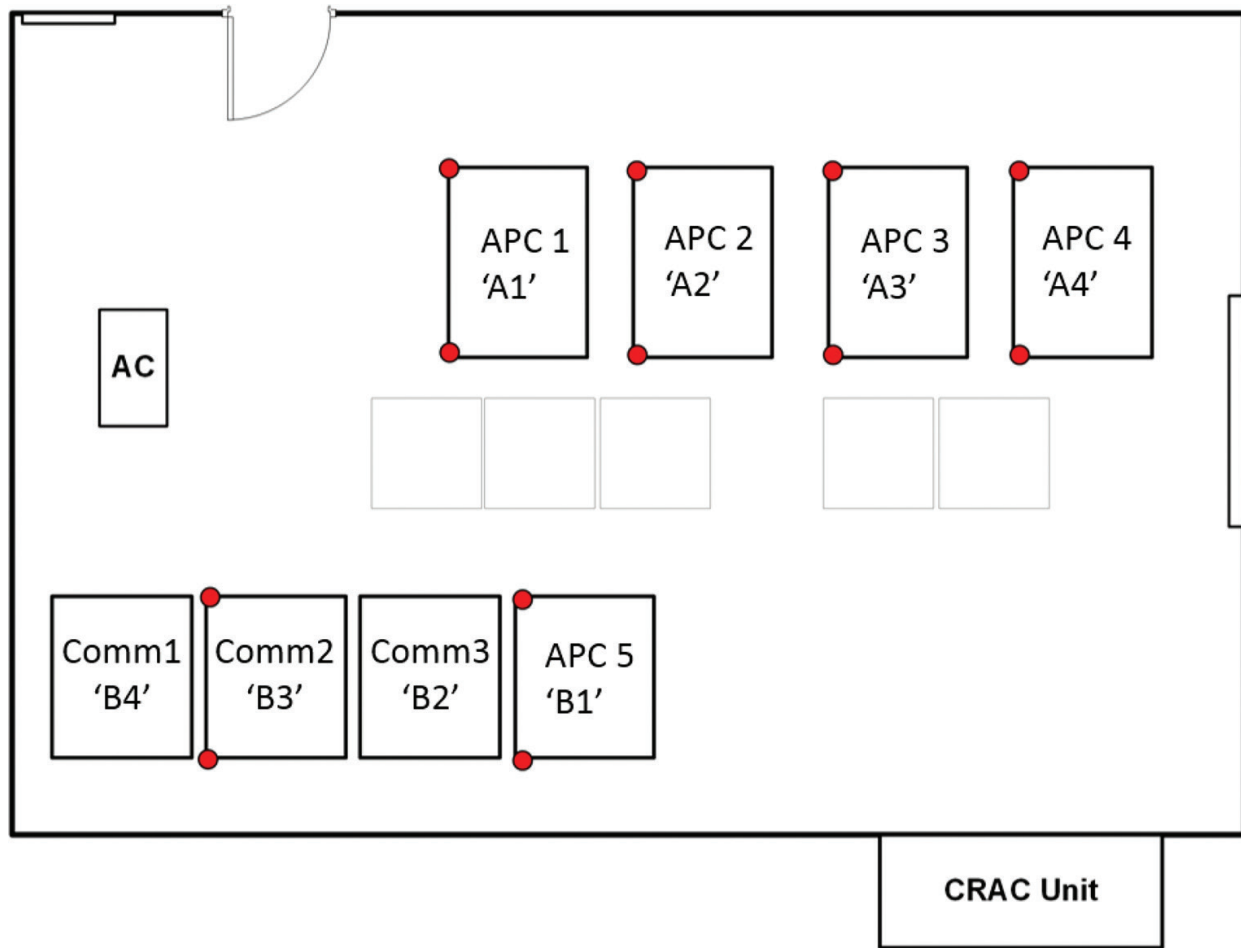


Figure 7. Floor plan of the data centre room used for the simulation-based assessment.

In order to demonstrate the functionality and feasibility of this approach, the Simulator GC implements a virtual data centre model that is based on the actual GENiC demonstration site, the C130 data centre at Cork Institute of Technology. The data centre space is cooled by one main computer room air conditioning unit (CRAC) and one backup air conditioning unit (AC) as illustrated in the floor plan depicted in **Figure 7**.

5.2. IT equipment and DC whitespace characteristics

To emulate the server workload in the data centre, a set of virtual machine (VM) configurations and the VMs' resource utilization traces are required. The traces used for the evaluation example presented here have been collected from a typical corporate data centre production environment and reflect typical enterprise workloads seen in a private cloud environment. The traces comprise resource utilization data for 2400 different VMs hosted on 132 servers. The key parameters of these servers are summarized in **Table 2**. The last column shows the number of servers of each specific type. Each server's dynamic power consumption is modelled as

$$P_{server} = (P_{max} - P_{idle}) \times u + P_{idle}$$

where u is the CPU utilization, P_{\max} is the server's power consumption at full load (i.e. $u = 1.0$), and P_{idle} is the server's power consumption at idle state (i.e. $u = 0.0$). The total power consumption of the 132 servers is 24.5 kW if all servers operate at full load.

For the simulation-based evaluation example, each server has been mapped to a specific rack space in the simulated data centre. **Table 3** shows this mapping.

Type	CPU size [vcores]	CPU speed [MHz]	Mem. [GB]	Max. power [W]	Idle power [W]	# Servers
S1	8	3200	16	90	30	3
S2	8	3200	32	95	35	8
S3	8	3200	64	105	45	48
S4	12	2000	64	130	70	2
S5	12	2000	128	140	80	12
S6	12	2000	256	160	100	23
S7	24	2700	128	300	140	19
S8	32	2000	128	400	270	14
S9	32	2900	128	460	300	3

Table 2. Server parameters.

Rack	Servers (top to bottom)	ΣP_{\max}
B1	2 × S5, 6 × S3, 6 × S6, 6 × S8	4.3 kW
B2	No active equipment; patch panels only	0 kW
B3	10 × S3, 6 × S3, 3 × S6, 4 × S7, 2 × S8	4.2 kW
B4	No active equipment; patch panels only	0 kW
A1	2 × S4, 3 × S1, 8 × S3, 8 × S7, 2 × S5	4.1 kW
A2	4 × S3, 2 × S2, 4 × S5, 5 × S7, 2 × S8, 3 × S3	3.8 kW
A3	4 × S8, 4 × S6, 7 × S3, 4 × S5, 4 × S6	4.2 kW
A4	3 × S9, 6 × S6, 4 × S3, 6 × S2, 2 × S7	3.9 kW

Table 3. Mapping of servers to racks in the virtual data centre.

5.3. Cooling system characteristics

The environment of the data centre is maintained at temperatures between 18 and 27°C with a relative humidity of 30–60% as recommended by ASHRAE [20]. The CRAC unit ensures the required indoor climate. Supply air is distributed through a raised floor and goes to front side of IT devices through perforated tiles. Return air is drawn by the CRAC unit below the ceiling as shown in **Figure 8**.

The conditions of circulating air are controlled in the CRAC unit by a direct expansion system. A condenser coil of the direct expansion system is cooled by glycol, and heat is rejected to the external ambient environment in a roof-mounted dry cooler. The process and devices involved are depicted in **Figure 9**.

There is also an auxiliary floor standing air conditioning (AC) unit placed in the room, as shown in **Figure 10**.




Figure 8. Schematic of hot and cold aisle arrangements without containments.




Figure 9. Main cooling system.




Figure 10. Auxiliary air conditioning unit.

6. Simulation-based assessment of energy management

The simulation-based evaluation of the GENiC energy management (EM) platform tests the interaction of short-term (S-T) actuation and long-term (L-T) decision-making on the virtual C130 data centre test-bed that replicates the physical processes occurring in the real data centre facility. This interaction and the components involved are shown in Figure 11.

A key component in all evaluations reported in this paper (and shown in Figure 10 via the arrows between components) is the Communication Middleware GC, which provides the

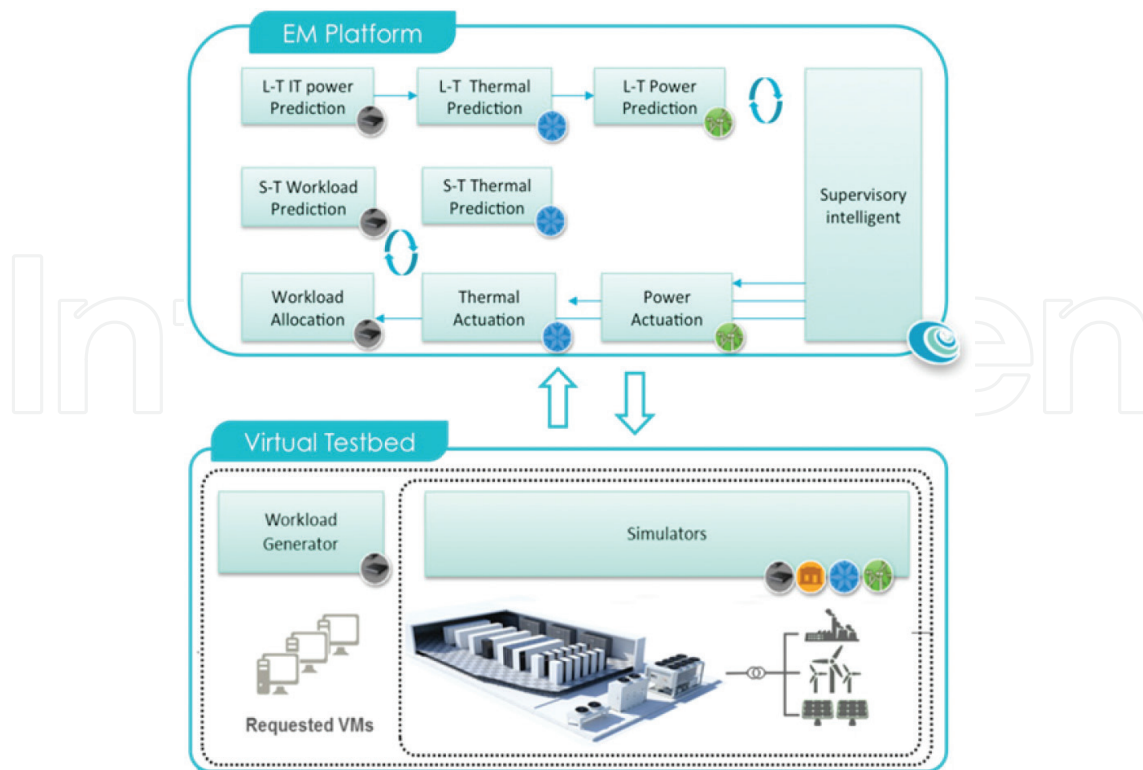


Figure 11. Interaction between EM platform GENiC components and virtual DC test-bed.

glue between all the different GENiC components and enables message exchange between components via the RabbitMQ broker (see above). The details of which components are relevant to a particular evaluation are discussed in the following.

6.1. Boundary conditions for the simulation-based assessment

All use cases are tested based on identical boundary conditions so that the different operating strategies can be compared to each other. The following external factors are considered as boundary conditions:

- **Requested VMs** are related to the type of services and end-user behaviour.
- **Electrical Grid info** is related to electricity market and the ratio of RES (CO₂ emission factor) in the grid.
- **Weather** conditions are specific to the DC location.
- **DC Operator Strategy** represents the baseline control strategy that establishes the reference baseline to assess the energy management saving potential.

6.2. Workload management GCG

The evaluation of the Workload Allocation Optimisation GC algorithms used within the GENiC prototype implementation was evaluated under the following scenarios (experiments):

- Workload Allocation—VM migration limits
- Workload Allocation—Thermal preferences

The experiment with VM migration limits refers to the evaluation of Workload Allocation Optimisation GC with different values for the maximum number of VM migrations allowed per time period. The evaluation with thermal preferences refers to the testing of Workload Allocation Optimisation GC considering a static thermal server preference when performing server consolidation. This experiment represents a thermal-aware workload allocation strategy [21]. The workload allocation experiment assesses the performance of the Workload Allocation Optimisation GC when it considers thermal actuation preferences. For the simulation-based evaluation, a static thermal preference matrix for each of the servers was developed based on Supply Heat Index (SHI) analysis [22] of the C130 data centre white space from the baseline inputs.

These scenarios were compared against each other and against a baseline allocation strategy. This comparison is assessed based on (i) the thermal behaviour in the white space (e.g. temperature distribution, hot spots) and (ii) energy consumption

6.2.1. GENiC components involved and testing process

The GENiC components involved in this particular workload management evaluation example are a subset of those that form the overall Workload Management GCG. This particular subset was chosen here to demonstrate the feasibility of the approach and demonstrate the overall system in operation. The experiments for this evaluation follow these steps:

1. The Simulators GC publishes the virtual time that synchronises the actions of the components involved in the experiment.
2. The Workload Generator GC publishes the VMs' resource utilization monitoring data for the current time step.
3. The Workload Allocation Optimisation GC optimizes the allocation strategy for the given arrangement in the virtual C130 DC.
4. The Workload Allocation Optimisation GC is able to consider thermal priority for each thermal box (where each thermal box represents one third of a rack). Static thermal priority is used to test a thermal awareness-based workload allocation strategy.
5. The Server Configuration component translates VM allocation to power consumption per box (one third of a rack).

The Simulators GC captures all the data relevant to this process for analysis and post-processing. The focus of this evaluation is to analyse the influence of workload allocation strategies on the temperature distribution of the white space as well as on the total DC energy consumption.

6.3. Thermal management GCG

Further experiments target the evaluation of the Thermal Management GCG algorithms with optimal thermal actuation. In this scenario, the GENiC prototype implementation is evaluated against a baseline operation strategy. This comparison is assessed based on data centre energy consumption and white space temperature distribution.

6.3.1. GENiC components involved and testing process

The GCs involved in this thermal management evaluation are a subset of those that form the Thermal Management GCGs. The subset chosen aligns with the requirements of the particular data centre demonstration site, and other, larger data centre configurations may use a broader spectrum of functionality. The experiments for the thermal management evaluation follow these steps:

1. Virtual synchronization time and current white space temperatures are published for the given time step.
2. The short-term (S-T) thermal prediction component predicts the thermal state of the white space for the next hour. This prediction supports the decision-making process that takes place in the Thermal Actuation GC.
3. Optimal temperature set points for the CRAC and AC units for the next time step are sent back to the HVAC systems model, which is part of the Simulators GC.

The Simulators GC captures all the data relevant to this process for analysis and post-processing. The focus of this evaluation is to analyse the influence of S-T prediction and thermal actuation strategies developed in the project on the temperature distribution of the white space as well as on the total DC energy consumption.

6.4. Power management GCG

In order to evaluate the power management aspects of the GENiC prototype platform, experiments were executed to evaluate the Power Management GCG algorithms under the following scenarios: (i) Power Actuation Logic, and (ii) Power Actuation Logic + SI static constraints. These scenarios are compared against each other and against the baseline operation. This comparison is assessed based on energy demand versus supply (broken down per source).

6.4.1. GENiC components involved and testing process

The GCs involved in this power management evaluation are a subset of those that form the Power Management GCGs and are selected to reflect the specific situation prevalent in the demonstration site. Elements of the power systems micro-grid available to the project, including a battery bank and an Organic Rankine Cycle (ORC), were modelled and included in this evaluation. The experiments for the Power Management evaluation follow these steps:

1. The Simulators GC generates the virtual time stamp and the current status of power metering for all equipment at the demand-side (DC) and at the supply-side (on-site RES).
2. The Power Actuation GC generates optimal set points for the batteries and the ORC plant for the next time step.
3. The Power Actuation GC receives a power policy (24 h profile) from the Supervisory Intelligence GC. A static SI constraint was used for the testing.

The Simulators GC captures all the data relevant to this process for analysis and post-processing. The focus of these experiments is it to analyse the power actuation operation strategies to satisfy the total DC demand. The power actuation real-time adjustments are defined so as to assure the renewable energy supply contribution. This is achieved through balancing the lack or excess of weather-dependent generation by using a controllable unit characterized with “unlimited” energy (kWh) capacity, which in this case is the ORC. The ORC has an unlimited energy capacity if the biomass storage is continuously refilled. It has to be understood that electrical batteries are characterised by limited energy capacity (here around 10 kWh) and limitations for the operation according to the definition of FSoC (fractional state of charge: between 0 and 1) upper and lower limits. According to the difference between weather-dependent renewable energy output prediction and real production, the ORC generation is adjusted taking into account the upper and lower power available referred to the maximum and minimum generation capacity of the ORC (here 4 kW minimum and 7 kW maximum).

7. Evaluation results

The simulation-based evaluation considers first results from the workload management experiments. The experimental set-up involved allocating workload over a 48-h period in a data centre using real VM resource utilization traces. Each VM was initially assigned to




Figure 12. Power consumption with different migration limits over 48-h horizon.

a particular server as per the real traces without the Workload Allocation Optimisation GC controlling the initial assignment. The only influence on power consumption was through VM migrations and server consolidation.

7.1. Workload allocation—VM migration limits

The first experiments evaluated the impact of the migration limit on the workload allocation (without thermal priorities for servers). This baseline is a migration limit of 0, that is, each VM was run on the server it was initially assigned to. Following from there, a series of experiments were executed to evaluate various migration limits (from 1 to 100) as shown in **Figure 12**.

As expected, increasing the migration limit resulted in a considerable reduction of power consumption (see **Figure 12**). The largest migration limit tested (100 migrations per 10 min time period) required just a few time periods to achieve a reduction from approximately 11 kW to just over 4 kW. Indeed, the average hourly energy consumption of the IT equipment was 6.71 kWh less with a migration limit of 100 than with the baseline. The figure for IT power consumption (see **Figure 12**) further illustrates that all positive migration limits tended to this equilibrium state, with a migration limit of 10 reaching the 4 kW mark in less than 9 h and the limit of 5 requiring approximately 24 h. Once reached, the variations in power consumption between the migration limits were minor. This means that if the workload allocator had controlled the initial assignment of VMs to servers, then a migration limit of 10 or even 5 would have been sufficient to achieve similar savings as with a limit of 100.

7.2. Workload allocation—thermal preferences

The experiments described in the following were performed under identical settings to those previously discussed with the exception that each server had an associated thermal preference, thereby allowing a proper ranking of servers. The thermal preference was used to rank the servers for consolidation.

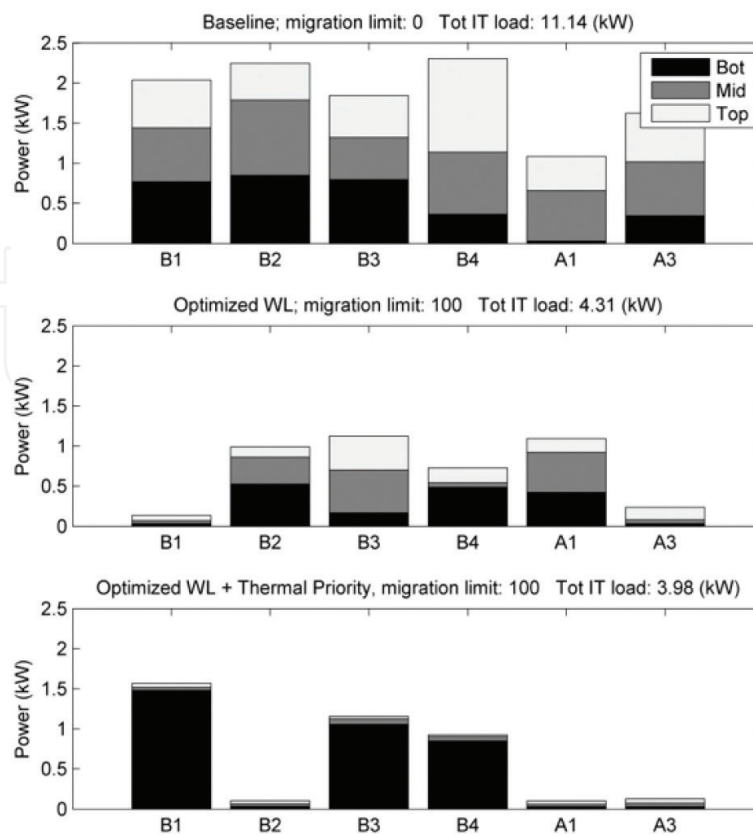


Figure 13. Workload distribution per third of rack.

In addition to the baseline described in the previous section, experiments were executed to assess power consumption with and without thermal preferences for migration limits of 10 and 100. The experiments showed that there is little difference in the total IT power consumption for the thermally ranked server consolidation, while HVAC energy consumption was reduced by approximately 20 kWh over the 48-h period relative to the baseline approach, and by 6.5 kWh compared to the scenario with 100 migrations and no thermal preference.

The behaviour of the scenarios with thermal preference can be better understood when analysed at the third of rack level (top, middle and bottom boxes) as shown in **Figure 13**. As can be observed, the only servers that were used by the GENiC energy management platform were those at the bottom level of three racks: B1, B3 and B4. The loads from all the other servers were migrated to servers in these locations and then servers that lost IT load were powered off, as can be seen from the power value for the scenario with thermal preference and limit of 100 migrations (bottom graph in **Figure 13**).

Finally, **Figure 14** presents the temperature distribution of the case study data centre C130 for (a) the thermal preference with 100 migrations and (b) the baseline. The baseline study indicates risks of a hot spot at the top layer of the last rack in row B. The supply air temperature is around 18°C; however, the inlet temperature of the particular box is approximately 23°C. The rise of temperature is due to infiltration of hot air from the hot aisle to the cold aisle space. The optimized workload allocation with thermal preference scenario ensures that the airflow

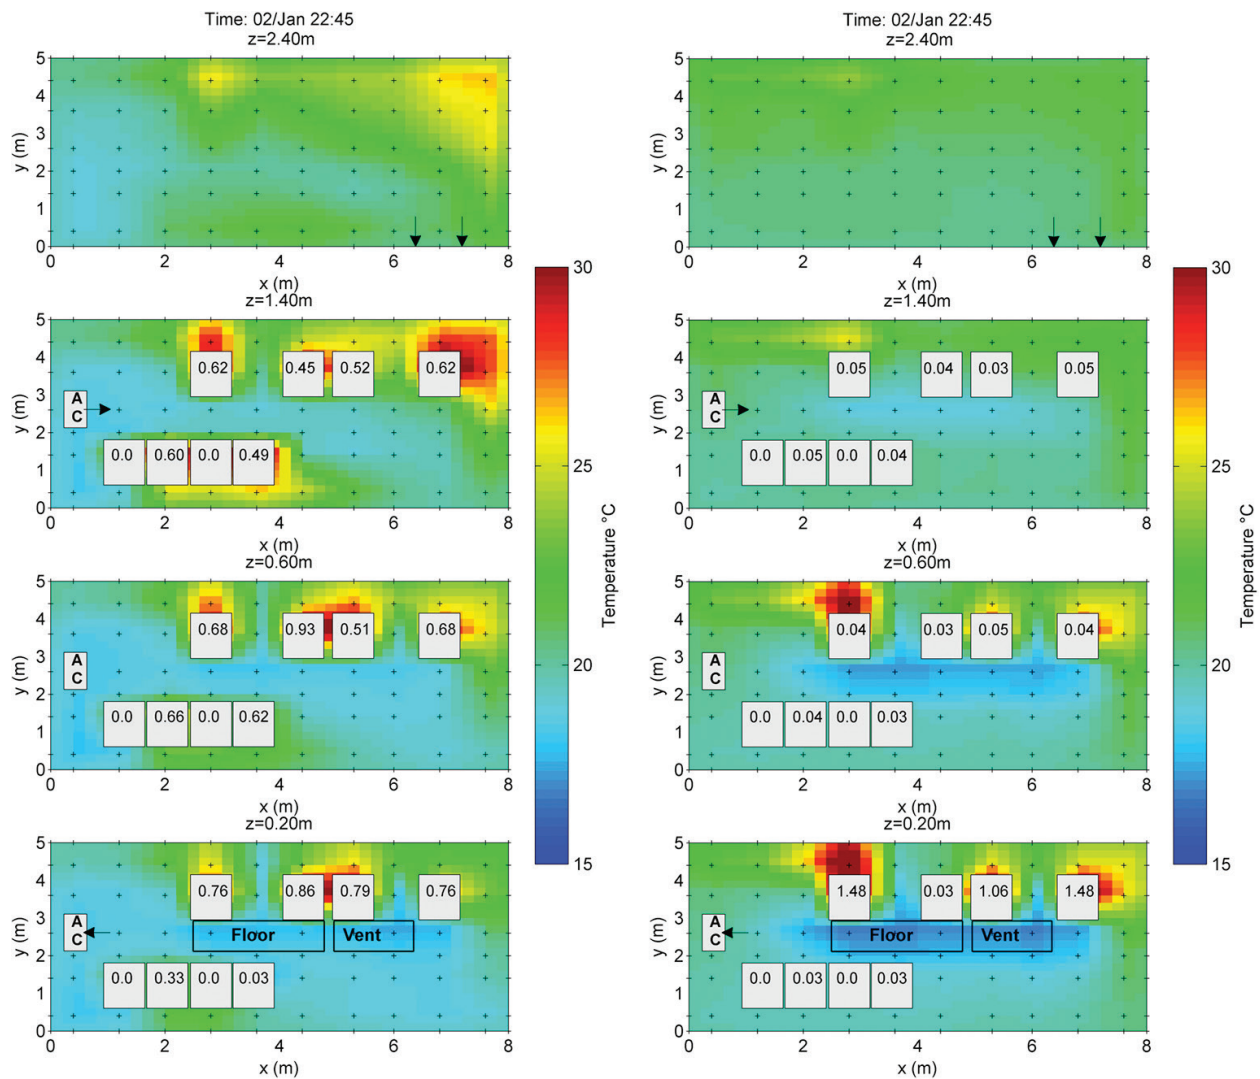


Figure 14. Temperature distribution for (a) thermal preference and (b) baseline.

will use the shortest path from the cold air supply to the heat source. The cold air is taken by preferable servers in the bottom boxes. The typical cold aisle-hot aisle distribution can be observed in this case. The inlet temperature of all active servers is approximately 18°C. This evaluation shows that the developed energy management platform can balance the temperature distribution in a data centre in such a manner as to avoid hot spots without the need for extensive structural changes to the cooling layout, for example, hot aisle containment.

8. Conclusions

In this chapter, an architecture for an integrated energy management system for data centres was presented. The architecture and prototype implementation was developed within the European Commission funded GENiC project. The proposed system combines optimisation of energy consumption by encapsulating monitoring and control of IT workload, data centre cooling, local power generation and waste heat recovery. The project conducted an initial

evaluation of the platform in terms of IT workload, thermal and power management based on a simulation model of a real data centre. The initial simulation-based assessment was chosen by the project for a number of reasons. It allows evaluating the performance of management and control algorithms before deployment in the real data centre space. Secondly, the architecture of the platform is designed such that the system interacts with the simulated data centre in the same manner as it interacts with the components in a real data centre, allowing also the testing and commissioning of novel management and control concepts before deployment in target space. The specific algorithms developed in the GENiC project attempt to optimise strategies focused on workload, thermal and power management in a data centre. The optimisation occurs at different time horizons, short-term predictions are generated to support actuation decisions that are made within each of the mentioned management groups, and long-term predictions supporting decision-making at the supervisory level (coordinating management groups). The evaluation presented in this chapter focused on an initial analysis of workload and thermal management techniques. The operation strategies applied by the Workload Allocation Optimisation GC prove significant savings potential (of up to 40%) in terms of total energy consumption. This reduction is achieved through the optimization of the allocation strategy of Virtual Machines (VMs) while switching off unused servers. The performance of the Workload Allocation Optimisation GC shows a more effective utilization of the data centre with the same number of processed IT jobs. The GENiC project will replace the simulation environment by a real physical data centre for the final evaluation and demonstration of the developed management algorithms and strategies in a real-world setting.

Acknowledgements

The authors acknowledge the European Commission's 7th Framework Programme in part funding the work reported here under Grant No. 608826.

Author details

Dirk Pesch^{1*}, Susan Rea¹, J. Ignacio Torrens², Vojtech Zavrel², J.L.M. Hensen², Diarmuid Grimes³, Barry O'Sullivan³, Thomas Scherer⁴, Robert Birke⁴, Lydia Chen⁴, Ton Engbersen⁴, Lara Lopez⁵, Enric Pages⁵, Deepak Mehta⁶, Jacinta Townley⁶ and Vassilios Tsachouridis⁶

*Address all correspondence to: dirk.pesch@cit.ie

1 Nimbus Centre, Cork Institute of Technology, Cork, Ireland

2 Building Physics and Services, TU Eindhoven, Eindhoven, The Netherlands

3 Insight Centre for Data Analytics, University College Cork, Cork, Ireland

4 IBM Research – Zurich, Rüschlikon, Switzerland

5 ATOS Spain SA, Madrid, Spain

6 United Technologies Research Centre Ireland, Cork, Ireland

References

- [1] The Independent, "Global warming: Data centres to consume three times as much energy in next decade", <http://www.independent.co.uk/environment/global-warming-data-centres-to-consume-three-times-as-much-energy-in-next-decade-experts-warn-a6830086.html>
- [2] Greenpeace, "How Dirty is Your Data? A Look at the Energy Choices that Power Cloud Computing. Greenpeace Report", <http://www.greenpeace.org/international/en/publications/reports/How-dirty-is-your-data/>, May 2011
- [3] Gao, P. X., Curtis, A. R., Wong, B., Keshav S. 2012. It's Not Easy Being Green. *Proceedings of ACM SIGCOMM*, Helsinki, Finland, pp. 211–222.
- [4] Qurush, A. 2010. Power-demand routing in massive geo-distributed systems, Ph.D. thesis, MIT.
- [5] NRDC Issue Briefing, "America's Data Centers Are Wasting Huge Amounts of Energy", IB:14-08-a, August 2014, <https://www.nrdc.org/sites/default/files/data-center-efficiency-assessment-IB.tif>
- [6] Uptime Institute, 2011. <http://uptimeinstitute.com>
- [7] <http://blogs.wsj.com/cio/2016/01/25/facebooks-irish-data-center-to-use-open-compute-hardware-renewable-energy>
- [8] <https://www.google.com/about/datacenters/renewable>
- [9] Deng, W., Liu, F., Jin, H., Li, B., Li, D. 2014. Harnessing renewable energy in cloud data centers: opportunities and challenges. *IEEE Network*, vol. 28, no. 1.
- [10] Cioara, T., Anghel, I., Antal, M., Crisan, S., Salomie, I. 2015. Data center optimization methodology to maximize the usage of locally produced renewable energy. *Sustainable Internet and ICT for Sustainability (SustainIT)*. Madrid, Spain, April 2015.
- [11] Das, R., Yarlanki, S., Hamann, H., Kephart, J. O., Lopez, V. 2011. A unified approach to coordinated energy-management in data centers. *Proceedings of the 7th International Conference on Network and Services Management (CNSM '11)*, pp. 504–508.
- [12] Jiang, T., Yu, L., Cao, Y. 2015. *Energy Management of Internet Data Centers in Smart Grid*. Springer Verlag, Berlin.
- [13] Pesch, D., et al. 2015. The GENiC Architecture for Integrated Data Centre Energy Management. in *1st Intl. Workshop on Sustainable Data Centres and Cloud Computing (in conjunction with IEEE UCC 2015)*. Cyprus, December.
- [14] GENIC. 2015. GENiC public deliverable D1.4 - Refined GENiC Architecture. <http://www.projectgenic.eu/>

- [15] Internet Engineering Task Force (IETF), "RFC7159: The JavaScript Object Notation (JSON) Data Interchange Format," March 2014, <https://tools.ietf.org/html/rfc7159>
- [16] RabbitMQ. 2016. <http://www.rabbitmq.com/>
- [17] A. I. Aravanis et al. 2015. "Metrics for Assessing Flexibility and Sustainability of Next Generation Data Centers," *2015 IEEE Globecom Workshops (GC Wkshps)*, San Diego, CA, pp. 1–6.
- [18] J. Ignacio Torrens et al. 2016. "Integrated Energy Efficient Data Centre Management for Green Cloud Computing - The FP7 GENiC Project Experience", in 6th International Conference on Cloud Computing and Services Science (CLOSER 2016), Rome, Italy, April.
- [19] Zavrel, V., Torrens Galdiz, J.I., Bynum, J.D. & Hensen, J.L.M. 2015. Model Development for Simulation Based Global Optimization of Energy Efficient Data Centres. *Building Simulation 2015: 14th International Conference of IBPSA*, December 7–9, 2015, Hyderabad, India (pp. 1079–1086). IBPSA
- [20] Tang, Q., Gupta, S.K.S., Varsamopoulos, G. 2007. Thermal-Aware Task Scheduling for Data Centers through Minimizing Heat Recirculation. *In Proceeding of IEEE International Conference on Cluster Computing (ICCC 2007)*. Austin, TX, USA, September 2007.
- [21] Sharma, R.K., Bash, C.E., Patel, C. 2002. Dimensionless Parameters for Evaluation of Thermal Design and Performance of Large-Scale Data Centers. *In 8th AIAA/ASME Joint Thermophysics and Heat Transfer Conference*. St. Louis, MO, USA, June 2002
- [22] ASHRAE. 2011. TC 9.9 Thermal Guidelines for Data Processing Environments – Expanded Data Center Classes and Usage Guidance. *Data Processing*: 1–45.

