4

Advanced Technical Concepts for Efficient Electrical Distribution and IT Management

Eduard Oró[1], Mauro Canuto[2] and Albert Garcia[1]

[1]Catalonia Institute for Energy Research – IREC, Spain
[2]Barcelona Supercomputing Center – Centro Nacional
de Supercomputación, Spain

4.1 Advanced Technical Concepts for Efficient IT Management

In this section, advanced technical concepts for efficient IT management are presented which aim in reducing the electricity demand of the IT equipment due to higher performance ratios by the use of consolidating tasks. Additionally, strategies for workload shifting are addressed.

Green algorithms rely on virtualisation [1] as a core technology that simultaneously executes full operating systems as guests within a single hardware node. Virtualisation brings the following advantages for both clients and cloud providers:

1. Physical hosts can be shared transparently to the clients, which are isolated as if each client were using a dedicated physical host. Virtualisation allows common users to get administrative permissions to configure the operating system, networking, and to install and uninstall software packages.
2. The resources (i.e. CPU or Memory) can be dynamically assigned and unassigned to the virtual machines (VMs) at runtime.
3. VM can easily migrate between physical resources at runtime. The migration is transparent to the user. Migration allows distributing VMs at runtime across the resources to increase server consolidation and save energy costs.

The flexibility of cloud computing and its success as a business model involve that users with diverse workload requirements access the cloud. Tasks handled by clouds may be CPU-intensive, I/O-intensive, memory-intensive, disk-intensive or even a combination of them. The high costs of energy for current Data Centres require, in addition to an efficient building and hardware design, the addition of policies and models to assess the VM allocation and management process to minimise the energy costs also from the software perspective.

Modern approaches for energy-aware IT management rely on modelling of the power consumption of the current workloads that run within modern hardware architectures [2]. There is a noticeable difference in power consumption when the tasks dominate different resources such as CPU, memory, network and hard drive [3]. Therefore, this means that a 100% HPC load which dominates mainly CPU is not the same that a 100% data load which dominates mainly hard drive. Chen et al. [4] build a linear power model that represents the behaviour of hardware nodes running VM under high-performance computing workloads. Their model may not be suitable to perform accurate predictions since it relies in hardware nodes. Several authors [5, 6] build power models to infer power consumption that apply to VM's power metering by using existing instrumentation in server hardware and hypervisors. However, none of these approaches considers the impact of hardware heterogeneity. In addition to the hardware heterogeneity, the heterogeneity of workloads (high-performance computing, data-intensive, real-time web workloads, etc.) must be also taken into account. The energy and power models of both hardware and workloads are the basis for cloud resources allocation and management algorithms. Hypervisor-level resource management methods (VM placement, resizing and migration) may be used to improve energy saving of cloud Data Centres [7, 8]. Consolidating the maximum number of VMs within the minimum number of physical hosts while turning off the idle hosts would minimise the energy impact while maximising the energy efficiency of the Data Centres. Thus, here it can be differentiated between two IT management strategies:

- **Consolidation strategy**. The objective of this strategy is to allocate the VMs, necessaries to satisfy the IT workload demand, in the minimum number of servers. Therefore, some servers are working at full load and the rest are kept in an idle state.
- **Turn-off idle servers**. It is a complementary strategy to the consolidation method, where the servers that are not being used are turned off, instead of being in an idle state.

As an example, Figure 4.1 shows the IT load of 10 different servers for two different scenarios. The reference scenario shows random allocation of IT load in all the available servers while the consolidation scenario concentrates the overall IT load only in the minimum required servers which in this case is 5 servers at full load. In this situation, the other servers (from 6 to 10) are in idle mode. Notice that when servers are in idle mode, they still consume energy and will depend on the typology and architecture of the servers. As an example, Figure 4.2 shows the nominal energy consumption for the same servers. Due to the implementation of consolidating strategies, Data Centre can save energy but will depend on the IT workload and the idle server energy consumption if this strategy will have impact or not.

Moreover, there is the possibility to turn off the servers in idle position. Figure 4.3 shows the nominal energy consumption in this situation, and it is clear that now the potential energy reduction is higher. Notice that the numbers are only to show the concepts and, in the reality, they can change in function of the IT load, servers configuration, etc.

However, these techniques would decrease the performance and the quality of service of the deployed services [9]. A trade-off is required between energy and performance that implies distributing (as opposite to consolidating) VMs with special performance requirements. Most works about consolidation of VMs focus on performance [10] and processor energy consumption [11] but do not consider the energy consumed by VM migration [12]. In the literature, it is

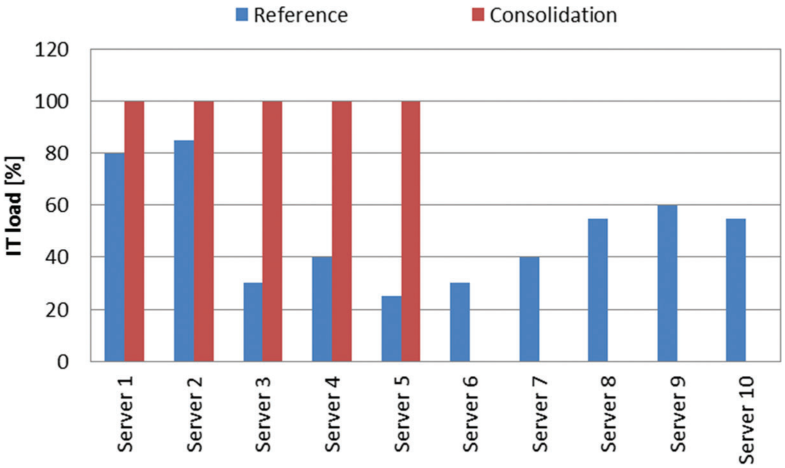

Figure 4.1 IT load for 10 servers with different IT management strategies (consolidation).

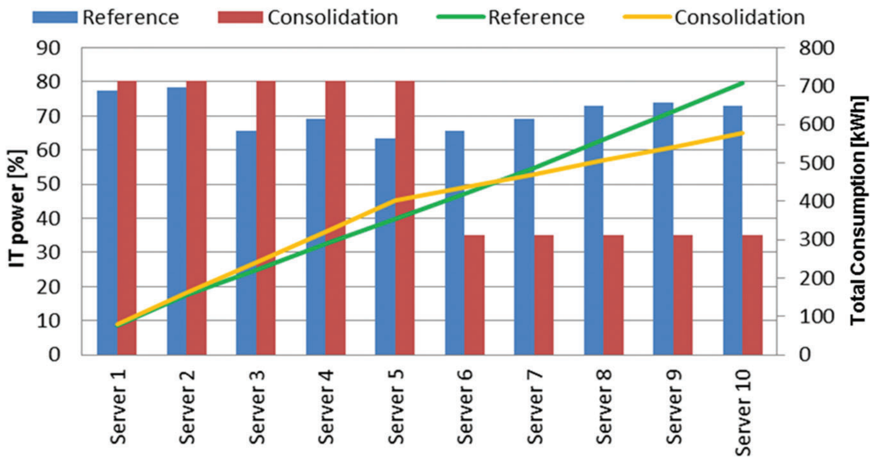

Figure 4.2 Nominal energy consumption for 10 servers with different IT management strategies (consolidation).

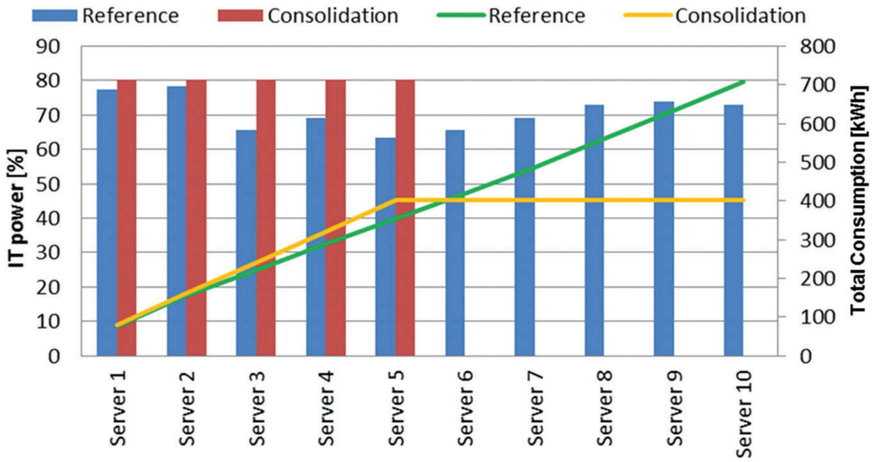

Figure 4.3 Nominal energy consumption for 10 servers with different IT management strategies (turn off servers at idle position).

available some work with potential energy savings up to 85% (but increasing the leeway by considering off-peak hours and bigger Data Centres) [13, 14]. Other approaches that are more realistic reported energy savings between 15 and 30% [15, 16]. Goiri et al. [15] measured an improvement of up to 15% when providing consolidation techniques that are aware of the power efficiency with the support for migration of VMs (Figure 4.4). The same authors showed how exploiting heterogeneity of Data Centres can improve power efficiency by up to 30% (Figure 4.5) while keeping a high level of SLA enforcement. In addition to the logical overhead caused by the overloading of tasks, the interferences between the different tasks that run in the same pool of resources must be considered [17–21]. Such interference would cause an overhead that consumes extra power and reduces energy efficiency.

When measuring the energy consumption of tasks within the resources, the thermal impact of the tasks must be considered (because the cooling systems also consume energy). It is required to calculate the relation between temperature and dissipated power for each node in order to engage thermal-aware temperatures, such as playing with "hot" and "cold" tasks to control temperature [22] or placing "hot" tasks to "cold" or best-cooled processors [23].

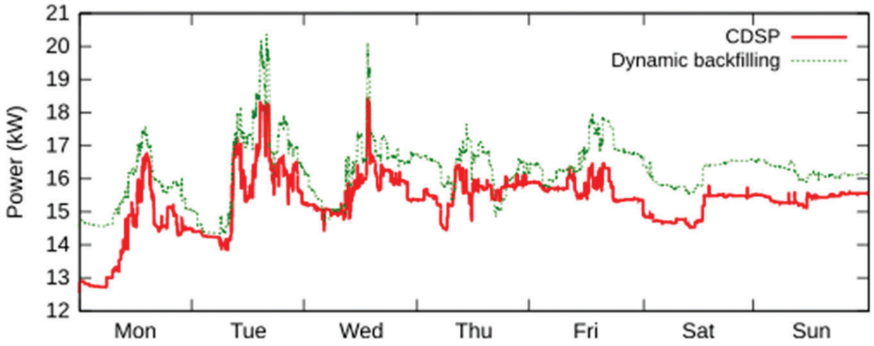

Figure 4.4 Power consumption comparison between consolidation policies (CDSP) and dynamic backfilling policy [15].

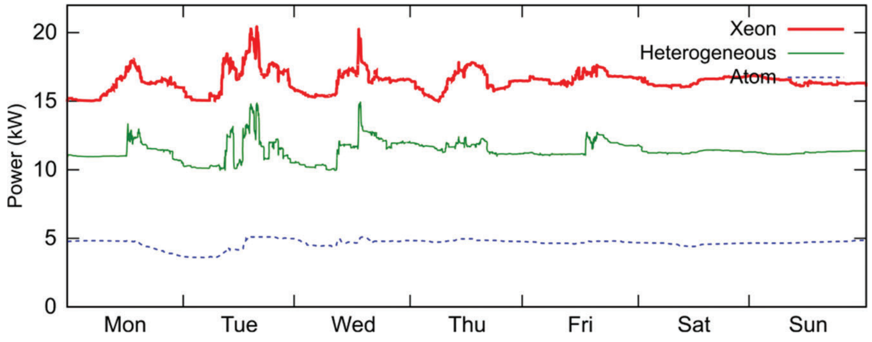

Figure 4.5 Power consumption with different levels of heterogeneity [15].

When the objective is not only to increase the energy efficiency but to maximise the ecological efficiency (reducing emissions and pollution), Data Centres that partially operate with renewable energies may schedule their workloads (if possible) according to the availability of such energies [24, 25]. For example, a Data Centre connected to solar PV schedules batch applications to the hours with the highest solar radiation. To show this, Figure 4.6 shows the IT power consumption for a Data Centre which is also connected to a PV system. Notice that during the peak PV production, the Data Centre is not consuming all the electricity generated by renewables and therefore should be sent to the main grid. On the other hand, when IT scheduling strategy is implemented in the Data Centre, the IT load can be adapted to the PV production as Figure 4.7 shows. In this situation, the Data Centre is fitted to the PV production increasing the share of renewables and decreasing the CO_2 emissions. Notice that in both scenarios, the IT consumption is the same but the execution time is different. This strategy cannot be implemented in web workload when the task should be done as a request for the user but for HPC and Data workloads exist the possibility to schedule the execution of task when it is more profitable for the Data Centre. As a complementary policy, Data Centres equipped with renewable energies that could be activated on demand (e.g. CHP or fuel cells running with biomass or biogas) could adapt energy supply to the expected workload [24].

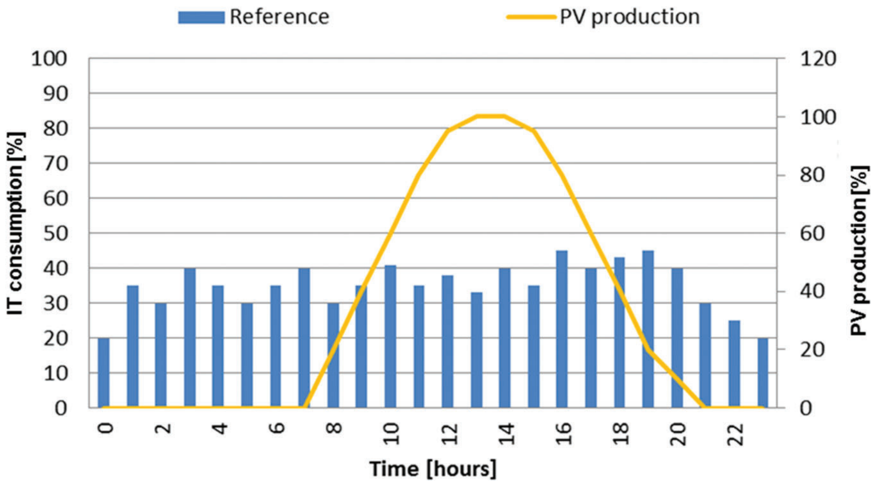

Figure 4.6 Schematic diagram of IT power consumption and PV power production without an optimisation.

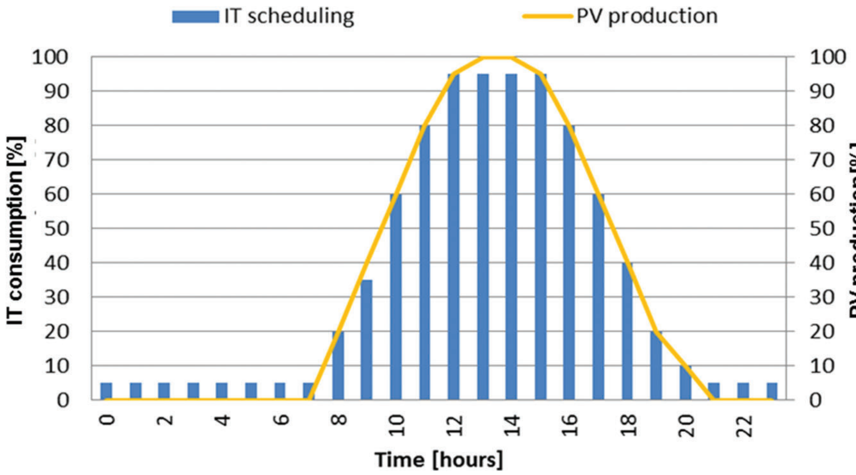

Figure 4.7 Schematic diagram of IT power consumption and PV power production under IT scheduling strategy.

Within the RenewIT project, a virtual machine manager (VMM) for Data Centre workload management has been developed to optimise the VMs placement in order to maximise the efficient usage of the resources (i.e. renewable energy) while keeping the workload performance standards that are mandatory by the cloud users. The VMM can be downloaded from the official BSC-RenewIT open source repository [26]. The main functionality of the VMM is to perform the deployment of VMs according to a policy specified by the owner of the infrastructure. Several policies have been included in the VMM:

- Distribution: It distributes the VMs trying to maximise the number of servers used. When using this algorithm, if two scenarios use the same number of servers, the one where the load of each server is more balanced is considered first. This policy is not energy saving but aims to maximise the performance and hence can be used for comparison purposes with the other policies.
- Consolidation: It distributes the VMs in a way that the number of servers that are being used are minimised.
- Energy aware: It deploys the VMs in the hosts where they consume less energy, according to the models and forecasts that were developed and released in Deliverable 2.2 of the RenewIT project [26].

In order to apply the scheduling policies described (distribution, consolidation and energy aware), the VMM needs to interact with other components such as the Energy Modeller (provides forecasts related to the energy consumption that a VM would have on a particular host) and the Infrastructure Monitor (to know the load of each host of the cluster – Ganglia [27] and Zabbix [28] – monitoring systems are used). Apart from the scheduling of VMs using the policies described above, the VMM offers other functionalities:

- It can manage the life cycle of VMs. This means that, by using the API that the VMM provides, it is possible to reboot, shutdown, suspend, restart and destroy VMs.
- It is also possible to perform queries to know at any moment the state of each of the VMs deployed and retrieve information about them: CPUs, reserved RAM and disk, IP address, the host where they are deployed, their creation date and time, etc.
- The VMM can also be used to manage the images from which VMs are instantiated. Specifically, using the VMM, it is possible to retrieve the information of all the images that have been registered, delete them and upload new ones from public URLs or from a local URI accessible by the VMM.
- Using the VMM, it is possible to retrieve information about the hosts available in the cluster where it is operating. It is possible to check the capacity of the hosts in terms of CPUs, RAM and disk as well as their current load. The VMM is also able to check the power consumption of each of the hosts at any given time.
- Finally, the VMM can be used to calculate energy estimates. Given the characteristics of a VM or a set of VMs, the VMM is able to calculate what their power consumption would be.

The main scientific contribution of the VMM developed in RenewIT project is the implementation of a scoring policy that feeds a heuristic local search algorithm that is able to improve the allocation and management of VMs to minimise the usage of energy while maximising the usage of VMs. The green algorithms developed allow setting the two already mentioned techniques: self-adaptation to enable migration of VMs at runtime and time switching to postpone batch jobs to the future. To show the VMM operative, Figure 4.8 shows the IT power consumption for two strategies, the first one is a random allocation of the IT jobs while the other is the consolidation, so the VMM distributes the VM to minimise the number of servers. Using the consolidation strategy, the total energy consumption is being reduced up to 13%. However,

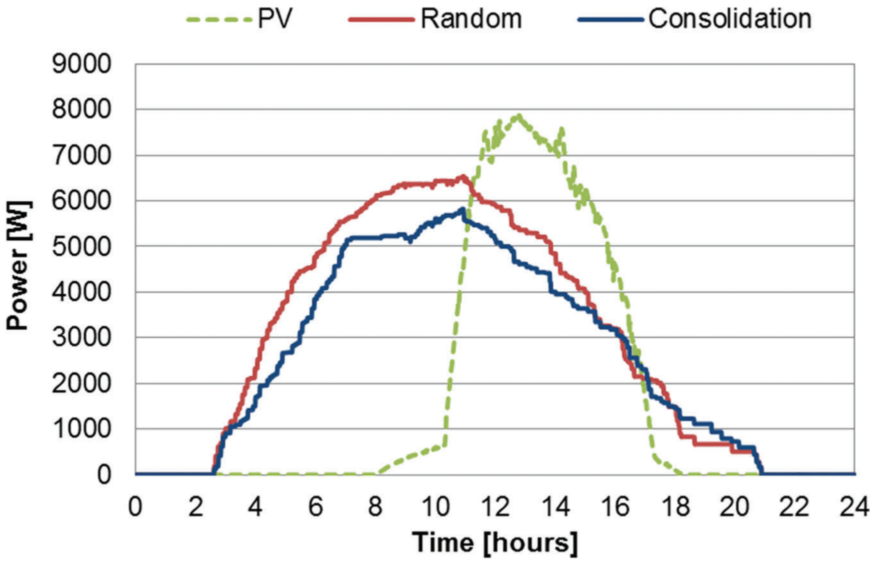

Figure 4.8 Random and consolidation strategies consumption for the same IT workload.

this strategy does not consider the availability of renewable energy sources which in this case is PV production. To greener the Data Centre, the energy-aware strategy is also implemented; here, the VMM locates the IT load to minimise the power consumption and maximise the usage of renewable energy of the Data Centre while fulfilling the performance requirements of all the VMs. Notice that there are restrictions on top of that to ensure that the tasks are always executed on time. Figure 4.9 shows the power profile of the power-aware strategy as well as the random strategy. It can be seen that the power aware is trying to follow the PV profile production over the day. In this situation, the energy consumption reduction is up to 25%, while the amount of green energy used from the PV has increased up to 60%, so the 60% of the total energy consumption under power aware is green energy.

The VMM demonstrated the feasibility of live migration for generic VMs in order to maximise the overall performance of the system in terms of energy efficiency. Moreover, to achieve the full utility of the self-adaptation policies, remote sleep/awake mechanisms must be provided to the physical nodes. This strategy will save energy when the consolidation policies allow complete freeing the load of a part of the resources.

When operating with different Data Centres that are geographically distributed, IT management policies could also minimise the energy impact and

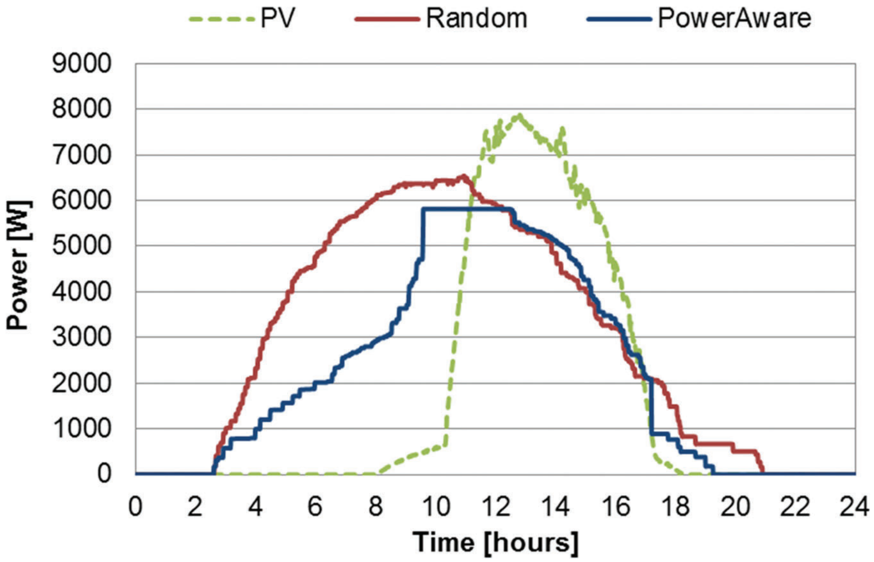

Figure 4.9 Random and power-aware strategies consumption for the same IT workload.

maximise the ecological efficiency by considering the spot status of each Data Centre [29–31]: energy cost, availability of energy according to their sources, or the status of the workloads and the physical nodes. Therefore, depending on the strategy, the Data Centre operators can manage the IT load between both infrastructures. As an example, Figure 4.10 shows the IT load in two Data Centres (A and B) as well as the percentage of renewables in the grid of each of the Data Centres. In this case, the objective is to run the IT workload using as maximum as possible the available renewable energy in the grid and therefore reducing the CO_2 emissions. Therefore, following this example is preferable to run the IT load in Data Centre A from 00:00 to 05:00 since the renewable energy ratio in the grid is higher in Grid A than in Grid B but after 05:00, it is better to move the IT load to Data Centre B. Therefore, when a geographically distributed IT strategy is implemented between Data Centre A and Data Centre B, the IT load is moved between the Data Centres, as Figure 4.11 shows. Notice that in both scenarios, the IT consumption is the same but the execution Data Centre is different.

Within the RenewIT project, also a VMM with a supra-Data Centre component has been developed that every time that a new VM is going to be deployed, it can choose the Data Centre destination in order to minimise the overall consumption or to maximise the use of on-site renewables or the

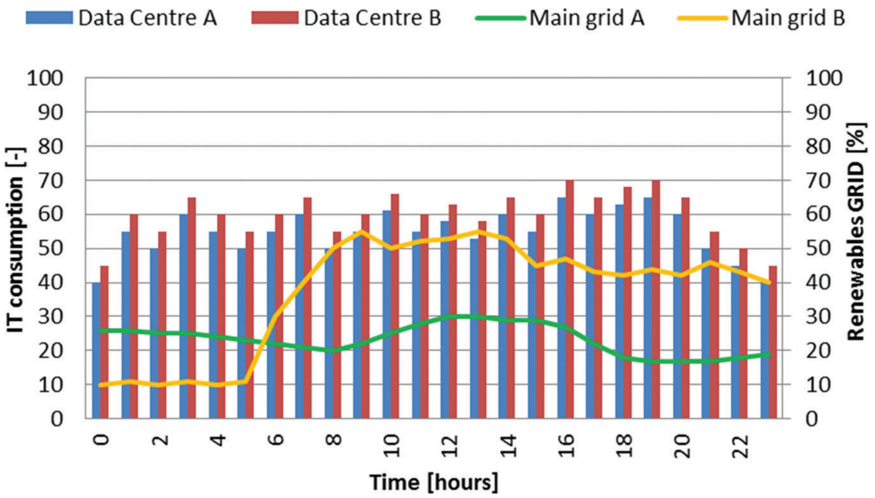

Figure 4.10 IT power consumption and percentage of renewables in the grid for 2 Data Centres.

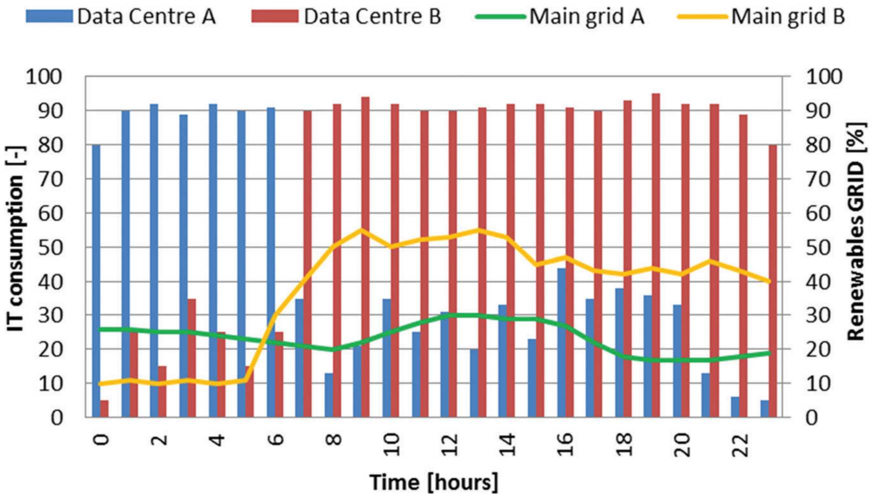

Figure 4.11 IT power consumption and percentage of renewables in the grid for 2 Data Centres under geographical IT scheduling strategy.

renewable grid ratio. As an example of the VMM deployment, Figure 4.12 shows the Data Centre power consumption of a random IT workload allocation strategy between two Data Centres. Notice that both facilities are located

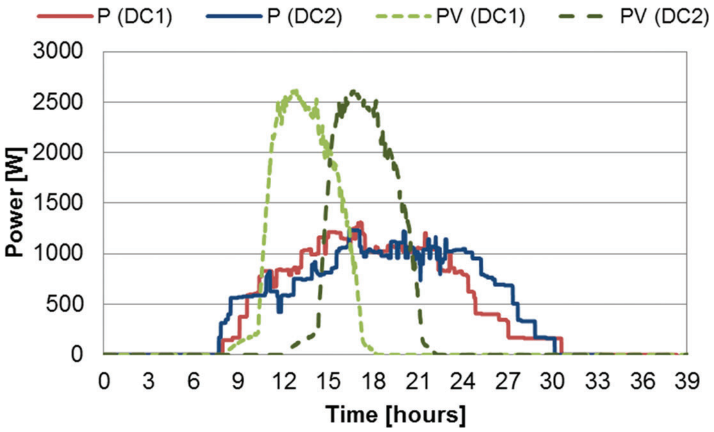

Figure 4.12 Energy consumption for an IT workload random allocation between two different Data Centres.

in different location as the PV power consumption shows (i.e. Athens and Seville). In this situation, the VMM does not considerate any relationship between the IT workload allocation and the green energy available in the facility, and therefore, the IT workload is allocated randomly in any of the two federated Data Centres. On the other hand, when the VMM is updated with the supra-Data Centre component, the VM will be deployed in one of the federated Data Centres to minimise the overall energy consumption or to maximise the use of on-site renewables. When this strategy is implemented in the same situation than explained before (Figure 4.13), the energy reduction is only being reduced by 2% but the share of renewables in the federated Data Centres is being increased up to 20%.

4.2 Advanced Technical Concepts for Efficient Electric Power Distribution

4.2.1 Introduction

The electrical distribution system of the Data Centre includes, among other equipment, switches, panels and distribution paths, UPSs and autonomous diesel generators for power backup in case of mains failure. These represent important expenditures for the Data Centre owner. However, since security of supply is principal for the installation, these expenditures are somehow unavoidable. The UPS is one of the main components of the electrical system of a Data Centre. This device is composed by an energy container (the most

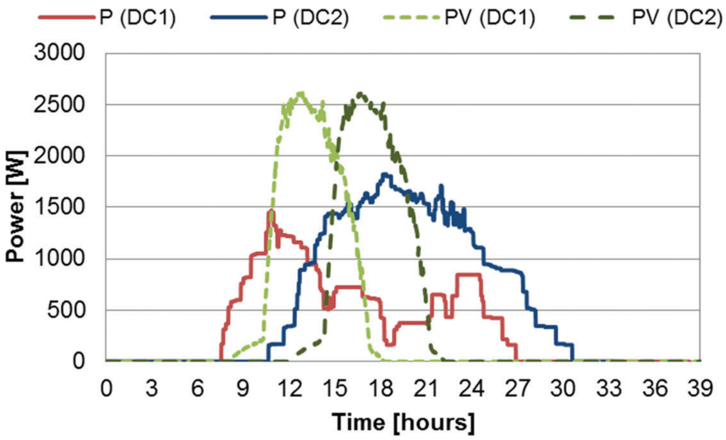

Figure 4.13 Energy consumption for an IT workload supra-Data Centre allocation between two federated Data Centres.

common is to use a battery bank or a flywheel) and a couple of power converters for its coupling to the electrical system of the facility. It is in charge of providing power to the IT equipment in case of electrical failure and before the electrical backup starts (actually a generator can operate at working conditions in less than a minute). UPS devices are one of the components of the electrical system that generates more losses. Usually, legacy UPS generates a 6% of losses compared with the total IT power consumption. Therefore, solutions able to minimise their losses have become an interesting topic in the Data Centre industry. As a result, a new generation of UPS, such as the modular or bypassed UPS, is emerging in the market. Another option for efficient electric power distribution is to enhance the energy capacity of the UPSs of the installation. This enhanced UPS could serve to provide energy during a few hours in case of a main failure as a diesel generator would provide or for peak hours. In this last scenario, the enhancement of the UPS would help to not oversize the main electrical elements to cover peak hours. However, the potential advantages of this advance electrical concept are related to the possibility of managing the energy stored in the UPS towards the technical and economic optimisation of the operation of the Data Centre. An example is to buy cheap electricity from the main grid for charging the electrical storages during the night and use it during the peak hours or when the electricity is too expensive. It is important to consider that for the economic suitability of installing energy storage systems, dedicated energy management algorithms are required. This gives the possibility to effectively time-shift the loads to low energy cost periods while still improving the power quality of the Data Centre and also being ready to act in case of a power failure in the main grid.

Therefore, in this section, a detailed behaviour of the advanced electrical concepts proposed in the framework of the RenewIT project is described. These concepts are as follows:

- Modular UPS
- Bypassed UPS in normal operating conditions
- Enhanced UPS for electrical energy storage.

4.2.2 Modular UPS

Typically, the capacity of the UPS exceeds the 100% of the maximum power of the installation and it is always connected, affecting the energy efficiency of the overall system. Here, a proposal to perform a modular design of the UPS is presented so that the number of modules connected in parallel can vary depending on the workload conditions. Each module can be activated or deactivated separately depending on workload to maximise the efficiency. The efficiency of a UPS module is particularly sensitive to workload conditions; it is minimal at low-load operation and increases until reaching its maximum at full-load operation. So adapting the number of connected modules to workload, i.e. adjusting the capacity of the UPS to the magnitude of the power to be transmitted, favours the operation of the system close to its ratings, thus maximising its efficiency. This concept is graphically depicted in Figure 4.14. Since the power is lower than 25% just one module is connected to the load while the other 3 modules are not. This obviously reduces the energy losses in those modules enhancing the overall energy efficiency.

The benefits achieved due to modularity are related with the increment of the load factor of the UPSs. The load factor is the percentage between the power demanded by the load and the nominal power of the UPS. Notice that not only the IT load profile but also relationship between the UPS design parameters and the IT installed capacity affects the instantaneous load factor. Moreover, the number of modules that the UPS will have affects the overall behaviour. After a literature research [32, 33] a UPS with 8 modules, when the UPS nominal power is between 50 and 800 kW, is used to demonstrate the feasibility of this concept.

The increment of UPS efficiency due to modularity needs to be related with the increment of the load factor between the UPS with and the one without modularity. Figure 4.15 shows the relationship between the load factor and

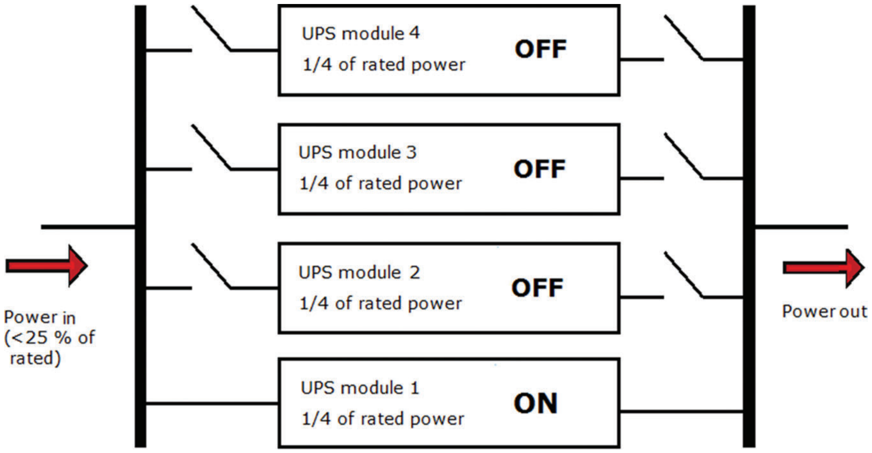

Figure 4.14 Graphical explanation for the concept of modular UPS.

the efficiency of a traditional and a modern UPS. Therefore, first it has to be calculated the overall load factor and then, using manufacturer efficiency curves, the overall energy efficiency of the UPS.

In the framework of the RenewIT project, the energy efficiency increase due to the use of modular UPSs has been done for different IT workloads (mainly HPC and web workload) and different IT power capacities.

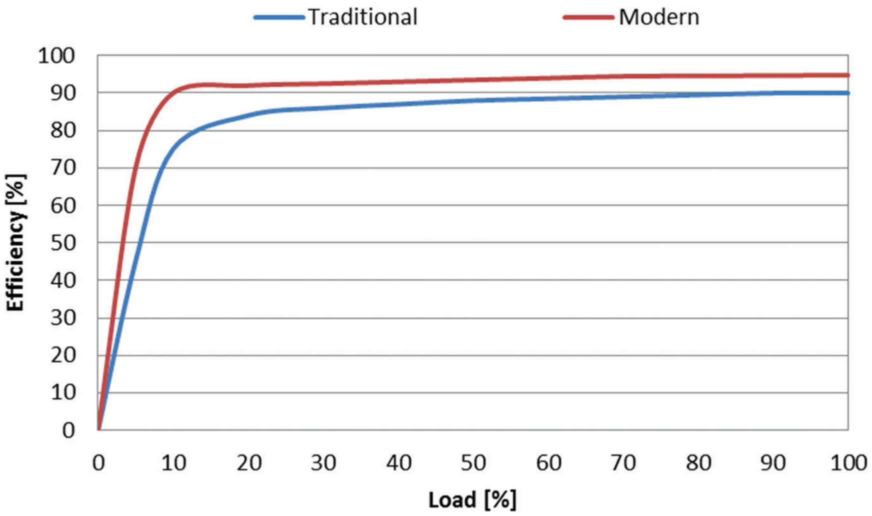

Figure 4.15 UPS efficiency depending on the load factor value.

As expected, modularity allows the UPS to work in a higher load factor and therefore increases its efficiency due to the fact that the system activates and deactivates UPS modules adjusting the nominal power close to the load demand. The results also show that the increment on the load factor in the scenarios with a web profile is lower than in the scenarios with a HPC profile. This is due to the variability of the web profile, which reaches lower load demands than in the HPC. Looking at the results in detail, the implementation of UPS modularity increases the average load factor of the UPS between 30 and 50%. This leads to an increment of 3–5% on the UPS efficiency, setting the average UPS efficiency around 97.5%. This improvement on efficiency means a reduction of 40–60% in the UPS losses, which results in 3–4% less power consumption in the overall Data Centre consumption.

4.2.3 Bypassed UPS

The majority of Data Centres utilises the static UPS in double conversion topology. In this topology, the critical loads of the Data Centre are always provided with fully conditioned mains power. Therefore, the alternating current (AC) currents and voltages from the main grid are converted to direct current (DC) and then back again to clean AC power. Applying this scheme, the losses in the UPS are noticeable because of the AC/DC double power conversion. Usually, legacy UPS generates 9% of losses compared with the total IT power consumption [34]. Therefore, solutions to minimise the UPS losses have become an interesting topic for Data Centre owners. Here, the bypassed UPS is described and analysed in detail. There are mainly three operational situations: standard situation or not bypassed, fully bypassed and partially bypassed, as Figure 4.16 shows. It can also be seen that the UPS avoids one or both converters depending on certain grid conditions and UPS characteristics, named partially and fully bypassed UPS, respectively.

Usually, UPS operates in normal conditions in order to provide the critical loads with “clean” AC power and ensure a constant power supply in case of mains failure, online operation results in no transfer time and the output voltage is of high quality. Underestimating these two main advantages, one can operate the UPS fully bypassed. In this operating mode, the power losses in the UPS are minimal since there is no power flow through the converters, but the critical loads are supplied with “dirty” power from the main grid. This reduction in power quality though can affect the operation and lifetime of the critical loads of the Data Centre and this is the main drawback of adopting this strategy. In order to improve power quality while still reducing

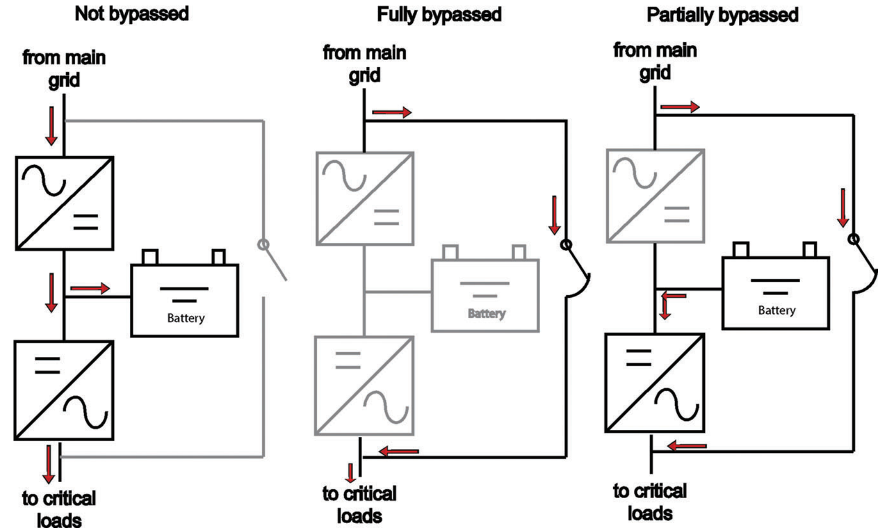

Figure 4.16 Operating modes for the UPS [35].

power losses in the UPS, it can be operated partially bypassed. In this configuration, the power flows from the main grid to the loads avoiding the AC/DC conversion in the grid side converter of the UPS; so this first power converter is offline. As a difference, the DC/AC power converter connecting the UPS to the loads is working, actively filtering the power from the main grid prior to reaching critical IT loads. As a result, this operation reduces power losses in the UPS while still improving the power quality for protection of critical loads of the Data Centre.

It is difficult to find a Data Centre operating at totally bypassed UPS due to the actual risk for this situation. Therefore, the increment of the energy efficiency of using only partially bypassed UPS in comparison with standard operational situation has been studied. The study is performed over different scenarios where different profiles and Data Centre redundancy levels are assumed. These profiles are web and HPC profile. The redundancy levels are as follows: redundancy level I, which means no redundancy in the electrical system, and redundancy level III, which means N+1 redundant component and two lines of electrical distribution, according to [36]. Moreover, in each scenario, IT power varies from 50 to 3,000 kW. On one hand, depending on the UPS topology, the transfer time between normal conditions and partially bypassed mode will be higher or lower. It is crucial to know the transfer time in order to ensure that the UPS will be able to provide power in case of main grid failures or clean the power in case of voltage disturbances, avoiding power interruptions in the IT equipment. A key factor during this process is the design ride-through of the IT power supply unit (PSU). PSU, in IT systems, can store small amounts of energy in capacitors and thereby allow a certain amount of ride-through time, which is the length of time that the IT equipment can continue to function during a complete loss of power. Usually, the ride-through time of a PSU ranges between 10 and 50 ms. Thus, it is imperative to know the ride-through capabilities of every PSU that will be powered by the UPS and know the transfer time of the UPS that needs to be lower than the PSU ride-through time. The green grid provides a summary table, for different UPS topologies, about the transfer time. This classification can be also found in the standard IEC 62040-3 [37]. On the other hand, to work with partially bypassed UPS, the grid quality also has to be considered. Even if nowadays some devices measure the quality of the power supplied by the grid, there is not too much information available in literature about the areas and the annual number of hours with an adequate grid power quality. Based on the specifications of some UPS manufacturer such as ENERGY STAR [38] and the information about grid power quality in the U.S., developed by the EPRI [39], a set of profiles showing the percentage of the annual hours working in partially bypassed mode are developed. Therefore, 3 scenarios have been considered, with the UPS working in partially bypassed mode at 25%, 75% and 95% of the annual hours.

The results of the study show that depending on the scenario (mainly IT workload and IT capacity installed) and the number of hours applying partially bypassed mode, the UPS efficiency can be improved between 0.5 and 3.0%. The highest UPS efficiency found was 98.2%, and the average UPS efficiency at partially bypass mode is around 97.5%.

4.2.4 Enhanced UPS for Electrical Energy Storage

The objective of this strategy is to enhance the capacity of the electrical energy storage in the Data Centre to optimise the energy flows into the facility. Most of the Data Centres have almost constant power consumption but in some cases, big differences can occur due to user's behaviour. In this situation, the enhancement of the electrical energy storage can really help to peak load shifting and therefore not only reduce the total capacity of the electrical elements and therefore reduce the investment cost but also to increase the energy efficiency of the system by working at higher load. Here, the enhancement of the electrical energy storage is aimed for electricity trade-off. That means if the electricity price is low, it is stored to be used when the electricity price is high. However, each scenario should be analysed in detail since the boundary conditions, in particular the electricity price, should change drastically between one and other Data Centre. In order to show the feasibility of this concept different IT capacities are studied (from 50 to 3,000 kW) and the following assumptions have been made:

- The UPS which are lead-acid batteries will only use the enhanced capacity to do smart trading. This enhancement can be up to 2.5 times the original capacity. They have a maximum nominal power of 709 kW, in function of the manufacturer specifications.
- Two different IT workload profiles have been analysed: web and HPC profile as Figure 4.17 shows.
- It is considered two scenarios for redundancy, a Tier I and a Tier III. On one hand, defining and enhancement of 100% in a 50 kW IT Data Centre with no redundancy (Tier I), the energy capacity of the UPS will be enhanced by 50 kW. On the other hand, for a 50 kW IT Data Centre with redundancy (Tier III), the total energy capacity of the UPS will be also enhanced by 50 kW, but there will be 2 UPS (line A and line B of the Data Centre, as shown in Figure 4.18.
- The electricity cost profiles are from Spanish industrial consumers (Figure 4.19).
- The economical evaluation is done in a 3-year scenario due to the fact that second life batteries should be replaced every three years of operation.
- It is assumed that the extra batteries are second life batteries, with an approximated cost of 150 €/kWh.

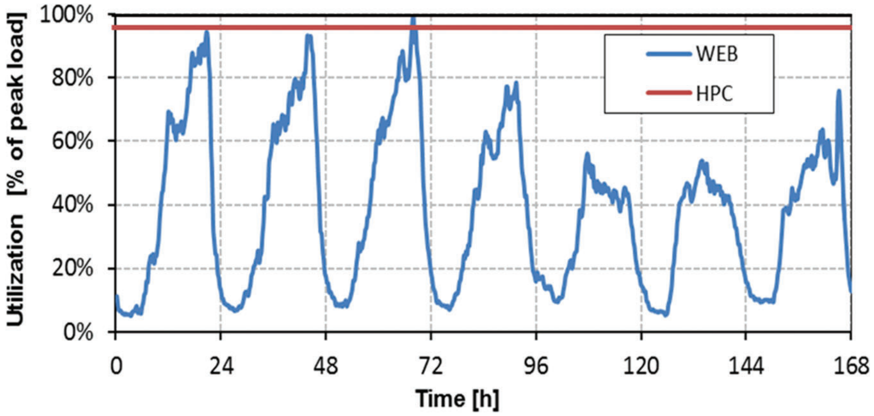

Figure 4.17 Utilisation defined by the IT load profiles web and HPC.

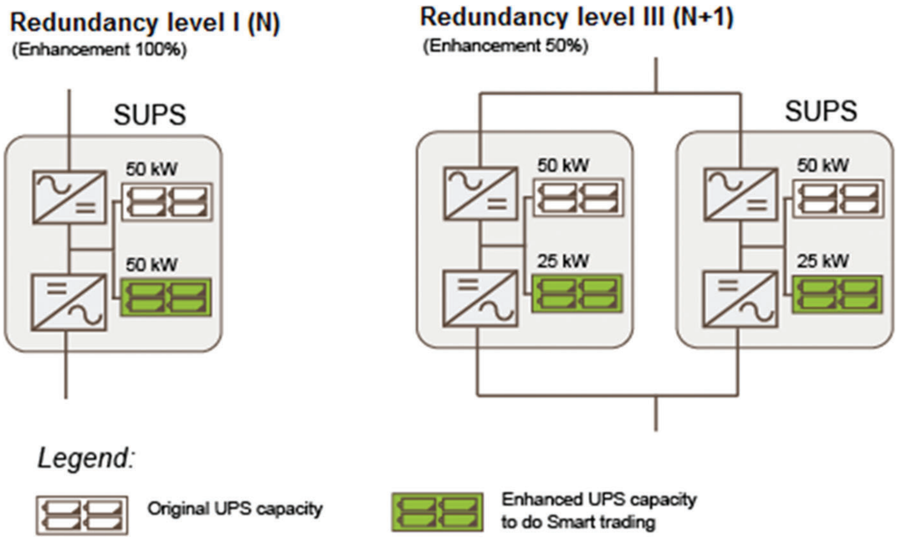

Figure 4.18 Enhancement of a UPS capacity to do smart trading, redundant and no redundant scenario [35].

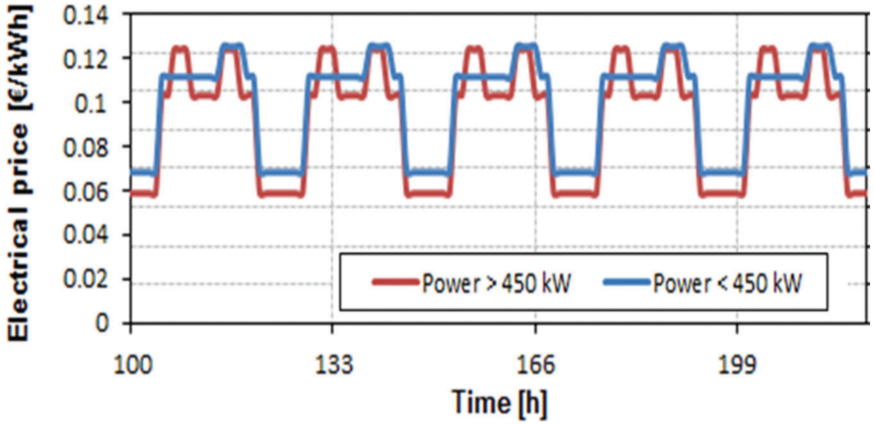

Figure 4.19 Electrical price profile for industrial consumers in Spain, five days of the year.

The smart trading algorithm is mainly governed by two conditions and three equations. The first condition compares the actual and the average electricity price. While the actual electricity price depends on the contract between the Data Centre and the electrical operator, the average is calculated using the electrical price of the last 24 hours. Depending on the comparative between the average and actual electrical price, the smart trading algorithm sends the order of charging or discharging to the UPS. If the actual price is higher than the average price, the smart trading algorithm will enforce the UPS to discharge the batteries. Otherwise, if the actual price is lower than the average, the smart trading algorithm will enforce the UPS to charge the batteries. The second condition is that the Data Centre energy management system cannot inject electricity (stored previously) to the grid. Therefore, during a discharge period, the state of charge imposed by the smart trading algorithm to the UPS will be restricted by the amount of energy that the UPS can deliver and the amount of power that the load is requiring. As example, if more power is demanded than the UPS can deliver, the state of charge becomes the minimal state of charge that the UPS can achieve, which is equivalent to discharge the extra batteries added at the UPS. Otherwise, if the UPS can deliver more power than demanded to cover the load, the smart trading algorithm calculates the exact state of charge necessary to satisfy the demand. Figure 4.20 represent the logical sequence followed by smart trading algorithm.

Where Equation (4.1) (minimal SOC), Equation (4.2) (Pmax,dch) and Equation (4.3) (reference SOC) are as follows:

$$\text{Minimal SOC} = \frac{100}{(\text{Enhancement} + 1)} \tag{4.1}$$

$$P\text{max}_{\text{dch}} = \frac{\left(\frac{100}{(\text{Enhancement}+1)} - \text{SOC}_{\text{bat}}\right) \cdot Kp \cdot \text{V}}{1000} \tag{4.2}$$

$$\text{MSOC}_{\text{ref}} = \frac{\text{PIT} \cdot 1000}{Kp \cdot \text{V}} + \text{SOC}_{\text{bat}} \tag{4.3}$$

The results of the study showed that with the boundary conditions used, especially the electricity price and the cost of the battery (150 €/kWh), there is no economic benefit of implementing extra energy storage in the system. These results do not mean that the concept is not working, it is just a matter of where to place the Data Centre and therefore which electricity cost is used, especially regarding the difference between day and night prices. In order to demonstrate this, it is calculated the maximum price of the second life battery in which the advanced energy efficiency strategy is starting to generate economic profits. For the current situation (Spain), this price is 73 €/kWh. Notice also that following the initial assumptions, not all the energy stored during a cheap electricity period could be later discharged when the prices are higher, in case that the battery could not be fully discharged before the control signal changes

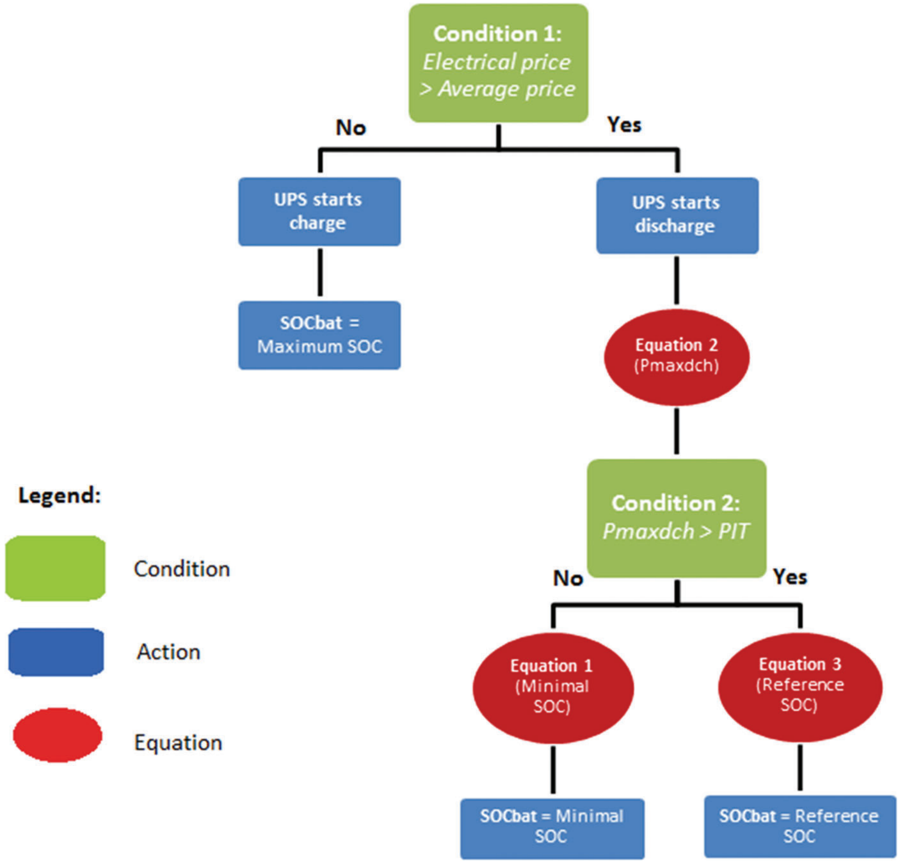

Figure 4.20 Block diagram of smart trading algorithm [35].

the UPS operational mode to charge. Additionally, when the UPS is storing energy, the IT load demand must also be satisfied. Thus, the amount of energy that can be stored by the batteries is limited because of the converter upstream of the battery. Therefore, the enhancement of the converter should also be taken into consideration for optimised strategies.

References

[1] Barham, P., Dragovic, B., Fraser, K., Hand, S., Harris, T., Ho, A., Neugebauer, R., Pratt, I., and Warfield, A. (2003). Xen and the art of virtualization. *ACM SIGOPS Operat. Syst. Rev.*, 37 (5), 164–177.

[2] Rivoire, S., Ranganathan, P., and Kozyakis, C. (2008). A comparison of high-level full-system power models, conference on Power–Aware computing and Systems (Hot Power '08)," p. 3.

[3] Smith, J. W., and Sommerville, I. (2011). "Workload classification and Software Energy Measurement for Rfficient Scheduling on Private Cloud Platforms," in *1st ACM Symposium on Cloud Computing (SOCC)*, October 2011.

[4] Chen, Q., Grosso, P., Van Der Veldt, K., De Laat, C., Hofman, R., and Bal, H. (2011). "Profiling energy consumption of VMs for green cloud computing," *IEEE 9th International Conference on Dependable, Autonomic and Sedure Computing (DASC)*, pp. 768–775, December 2011.

[5] Kansal, A., Zhao, F., Liu, J., Kothari, N., and Bhattacharya, A. (2011). "Virtual machine power metering and provisioning," *1st ACM Symposium on Cloud Computing*, pp. 39–50, October.

[6] Bohra, A., Chaudhary, V. and Meter, V. (2010). "Power modelling for virtualized clouds, *IEEE International Symposium on Parallel and Distribution Processing, Workshops and PhD Forum (IPDPSW)*, pp. 1–8, April 2010.

[7] Bo, L., Jianxin, L., Jinpeng, H., Tianyu, W., Qin, L., and Liang, Z. (2009). "EnaCloud: an energy-saving application live placement approach for cloud computing environments," *International Conference on Cloud Computing*.

[8] March, J. L., Sahuquillo, J., Petit, S., Hassan, H., and Duato, J. (2011). Real-time task migration with dynamic partitioning to reduce power consumption. *Actas de las XXII Jornadas de Paralelismo*, 185–190, September 2011.

[9] Ye, K., Huang, D., Jiang, X., Chen, H., and Wu, S. (2010). Virtual machine based energy-efficient data center architecture for cloud computing: A performance perspective. *International Conference on Green Computing and Communications and International Conference on Cyber, Physical and Social Computing*, 2010.

[10] Lu, L., Zhang, H., Smirni, E., Jiang, G., and Yoshihira, K. (2013). Predictive VM consolidation on multiple resources: beyond load balancing, *21st International Symposium on Quality of Service*, 2013.

[11] Deng, W., Liu, F., Jin, H., Liao, X., Liu, H., and Chen, L. (2012). Lifetime or energy: consolidating servers with reability control in virtualized cloud dataceneters, *4th International Conference on Cloud Computing Technology and Science*, 2012.

[12] Beloglazov, A., and Buyya, R. (2012). Optimal online deterministic algorithms and adaptive heuristics in Cloud data centers. *Concurrency Comput. Practice Exp.*, 24 (13), 1397–1420, September 2012.

[13] Uddin, M., and Rahman, A. A. (2010). Server consolidation: An approach to make data centers energy efficient & green. *Int. J. Sci. Eng. Res.*, 1 (1), October 2010.

[14] VMWare, "How VMware Virtualization Right-sizes IT Infrastructure to Reduce Power Consumption," 2011.

[15] Goiri, Í., Berral, J. L., Fitó, J. O., Julià, F., Nou, R., Guitart, J., Gavaldà, R., and Torres, J. (2012). Energy-efficient and multifaceted resource management for profit-driven virtualized data centers. *Future Generation Comput. Syst.*, 28 (5), 718–731.

[16] Beloglazov, A., and Buyya, R. (2014). OpenStack neat: A framework for dynamic and energy-efficient consolidation of virtual machines in OpenStack clouds. *Concurrency and Comput. Practice Exp.* [online].

[17] Mars, J., and Tang, L. (2013). Whare-map: heterogeneity in "homogeneous" ware house scale computers, *40th ACM/IEEE International Symposium on Computer Architecture (ISCA)*, pp. 619–630, June 2013.

[18] Mars, J., Tang, L., Hundt, R., Skadron, K., and Sofia, M. L. (2011). "Bubble-up – increasing utilization in modern warehouse scale computers via sensible co-locations," *44th IEEE/ACM Symposium on Microarchitecture (MICRO'44)*, pp. 248–259, 2011.

[19] Koh, Y., Knauerhase, R., Brett, P., Bowman, M., and Wen, Z. P. (2007). "Ana analysis of performance interference effects in virtual environments," *IEEE International Symposium on Performance Analysis of Systems and Software (ISPASS)*, pp. 200–209, April 2007.

[20] Somani, G., and Chaudhary, S. (2009). "Application performance isolation in virtualization," *IEEE International Conference on Cloud Computing (CLOUD'09)*, pp. 41–48, September 2009.

[21] Kousiouris, G., Cucinotta, T., and Varvarigou, T. (2011). The effects of scheduling, workload type and consolidation scenarios on virtual machine performance and their prediction through optimized artificial neural networks. *J. Syst. Software*, 84 (8), 1270–1290, August 2011.

[22] Yang, J., Zhou, X., Chrobak, M., Zhang, Y., and Jin, L. (2008). "Dynamic thermal management through task scheduling," *IEEE International Symposium on Performance Analysis of Systems and Software (ISPASS'08)*, pp. 190–201, April 2008.

[23] Vanderster, D. C., Baniasadi, A., and Dimopoulos, N. J. (2007). Exploiting task temperature profiling in temperature-aware task scheduling for computational clusters. *12th Asia – Pacific Conference on Advances in Computer Systems Architecture (ACSAC'07)*, pp. 175–185.

[24] Li, C., Zhou, R., and Li, T. (2013). Enabling distributed generation powered sustainable high performance data centre. *International Symposium on High-Performance Computer Archhitecture (HPCA)*, pp. 35–46, February.

[25] Goiri, I., Le, K., Haque, M. E., Beauchea, R., Nguyen, T. D., Guitart, J., Torres, J., and Bianchini, R. (2011). Greenslot: scheduling energy consumption in green datacenters. *International Conference for High Performance Computing, Networking, Storage and Analysis (SC'11), Article No. 20*, 2011.

[26] Macías, M., Canuto, M., Ortiz, D. and Guitart, J. (2015). Green management of data centres: model for energy and ecological efficiency assessment. RenewIT deliverable number 2.2.

[27] "Ganglia Monitoring System," [Online]. Available: ganglia.sourceforge.net

[28] "Zabbix: An enterprise-class open source distributed monitoring solution," [Online]. Available: www.zabbix.com

[29] Garg, S. K., Yeo, C. S., Anandasivam, A., and Buyya, R. (2011). Environment-conscious scheduling of HPC applications on distributed Cloud-oriented data centers. *J. Parallel Distributed Comput.* 71 (6), 732–749, June 2011.

[30] Pierson, J. M. (2011). Green task allocation: Taking into account the ecological impact of task allocation in clusters and clouds. *J. Green Eng.* 129–144, January 2011.

[31] Elmroth, E., Tordsson, J., Herandez, F., Ali-Eldin, A., Svard, P., Sedaghat, M., and Li, W. (2011). Self-management challenges for multi-cloud architectures. *4th European Conference on Towards a Service-Based Internet*, pp. 28–49, 2011.

[32] 2015. [Online]. Available: http://www.stat.osu.edu/∼comp_exp/jour.club/CamCarSal_EngModellingSoftware-2007.pdf

[33] Chloride Trienergy, Emerson, Modular architecture for high UPS Systems, White Paper, 2011.

[34] Rasmussen, N. (2007). Electrical efficiency modeling for data centres. White Paper #113, APC, Schneider, 2007.

[35] N. E. A. Shrestha, "Deliverable D4.5 Catalogue of advanced technical concepts for Net Zero Energy Data Centres. Final version," 2015.

[36] Uptime institute, “Data center site infrastructure Tier standard: Topology,” 2010.

[37] International Standard, IEC 62040-3, “Uninterruptible power systems (UPS), PART 3: Method of specifying the performance and test requirements, First edition, 1999–2003,” no. First edition, 1999–2003.

[38] ENERGY STAR, *Program Requirements for Uninterruptible Power Supplies (UPSs), Test Method final*, May 2012.

[39] EPRI, “Distribution System Power Quality Assessment: Phase II – Voltage Sag and Interruption Analysis,” Palo Alto CA, 2003.

[40] “OptaPlanner documentation,” [Online]. Available: http://www.optaplanner.org/learn/documentation.html