SEVENTH FRAMEWORK PROGRAMME

GreenDataNet

D3.3 - Electricity Consumption Forecasting Tool Design and Implementation

Final

[Manager]
Prof. David Atienza
Embedded Systems Laboratory (ESL) – EPFL

Rev 3.0

Contributors:

- Pablo García, EPFL.
- David Brunelli, Univ. Trento.

TABLE OF CONTENTS

REVISION SHEET

Revision Number	Date	Brief summary of changes
Rev 1.0	**09/01/2015**	Baseline document
Rev 2.0	**21/01/2015**	Revised structure and streamlined the experiments
Rev 3.0	**22/01/2015**	Index regenerated and final check
Rev 3.1	**25/01/2015**	Final check by EATON

1. INTRODUCTION

The main goal of the GreenDataNet project is to apply new technologies and tools to enhance the power efficiency of the urban data centres, effectively decreasing their average Power Usage Effectiveness (PUE), and making them more sustainable in the future.

To optimize the operation of a data centre, it is crucial to minimize both information technology (IT) and cooling energy consumptions. In particular, this deliverable presents the GreenDataNet Electricity Consumption Forecasting Tool (ECFT), a multi-level and multi-objective framework (composed by a combination of software models, and dedicated hardware) for the optimization of green virtualized data centres.

Figure 1.1 depicts a green data centre computing infrastructure (servers, storage, racks, etc.) having the IT equipment connected to the main Grid, and with PV modules and energy storage devices installed. The IT equipment and cooling system inside data centres operate at very low voltages, using much less power than the entire facility. To this end, the facilities are combined using higher voltage Power Distribution Units (PDU) that eventually connect to Uninterruptible Power Supplies (UPS) that serve the whole facility [3]. In this case, the renewable power is connected to UPS, as well the battery banks by means of the bidirectional Charge Transfer Interconnect bus (CTI-bus). Renewable energy and batteries are not intended as a protection from input power interruptions, but they must sustain long (i.e. hours) on-battery runtime.

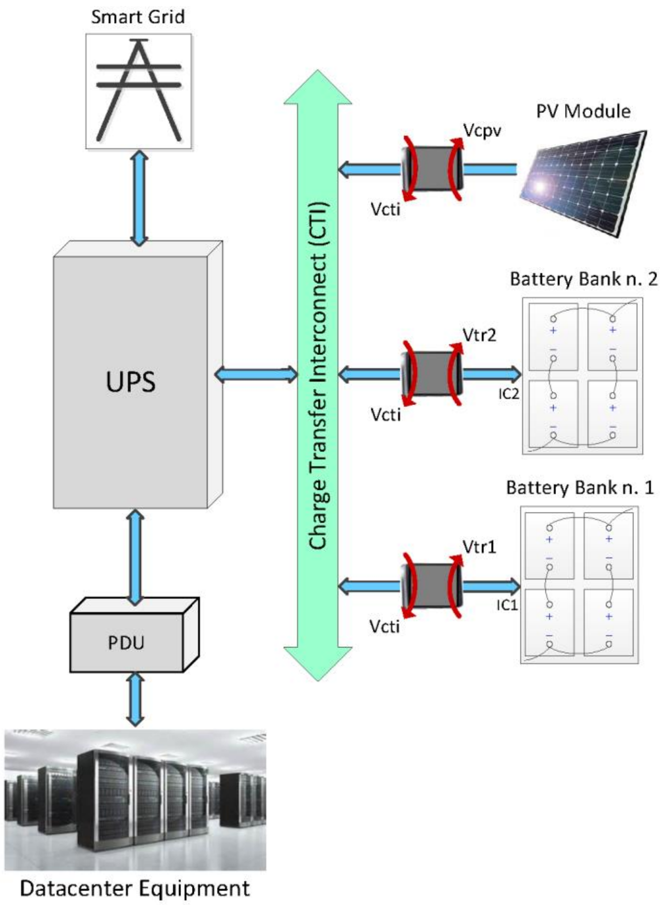

Figure 1.1 – Modelled green data centre structure (including only IT Infrastructure, but not the data centre building modelling and controller, described in Deliverable 3.7 – "Smart Energy Management System")

The two elements that compose the ECFT are the Green and the IT Infrastructure Energy Controllers. They are both described in this document, whereas additional details and performance evaluation are included into Deliverable 3.2 [1]. The overall diagram of our joint Green Energy and IT Infrastructure Energy Controllers is shown in Figure 1.2.

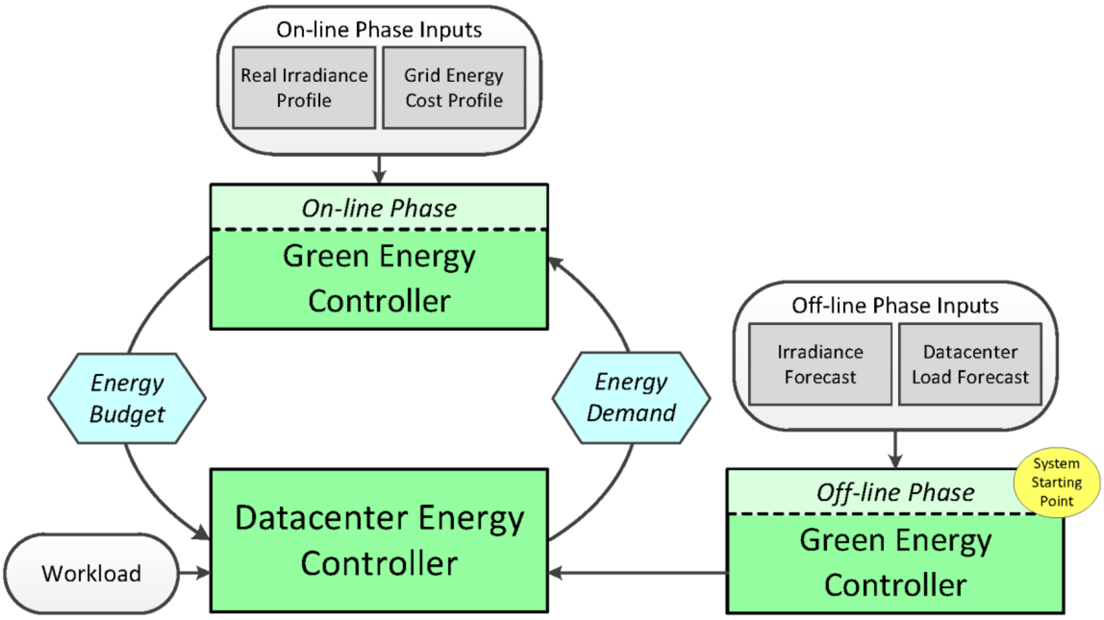

Figure 1.2 – Diagram of the Electricity Consumption Forecasting Tool, showing the Green and IT Infrastructure Energy Controllers (i.e., used as parts of the data centre building modelling and controller, which will be described in Deliverable 3.7 – "Smart Energy Management System"

At the beginning of the simulated time horizon (off-line phase), the Green Energy Controller computes the expected energy budget for the data centre using the information from the analytical models [2] (cf. see further details in deliverable 3.1), processing historical data centre power profiles as well as the sun irradiance forecasts. In this phase, only the primary battery bank is managed and no communication takes place, the goal is to obtain a rough schedule that maximizes the primary battery lifetime and minimizes the total energy cost.

The on-line phase starts when the Green Controller sends the available energy budget to the data centre Controller for the first time slot. From now on, the two controllers work concurrently exchanging messages, as depicted in Figure 1.2. After receiving the real energy consumption of the data centre, the Green Controller compensates the differences between: expected and available green energy (i) and IT Infrastructure Energy consumption against budget (ii) using the secondary battery as additional energy reserve. To this end, if the actual energy consumed by the data centre is higher than the expected, the Green Energy Controller compensates the IT Infrastructure Energy requirements. At the end of each time-slot the Green Controller provides an updated budget to the other. On the other side, the IT Infrastructure Energy Controller tries to find the best allocation for VMs on the servers using the VMs specification as incoming workload and the energy budget provided by green energy controller. The goal is to allocate VMs that yield in optimized total energy consumption of the data centre. Both controllers are invoked periodically. The following sections describe their design in detail.

1.1. DOCUMENT OVERVIEW

This deliverable details the GreenDataNet Electricity Consumption Forecasting Tool, that has already been presented in the introduction. The rest of the document is organized as follows: Sections 2 and 3 explain in detail the Green Energy and the IT Infrastructure Energy Controllers, respectively. Then, Section 4 describes how both controllers interact inside the Electricity Consumption Forecasting Tool (ECFT) and presents a set of experiments. Finally, Section 5 summarizes the main conclusions of this deliverable.

2. THE GREEN ENERGY CONTROLLER

The Green Energy Controller is an algorithm created to optimize the HES of Green Data centres with a HW support similar to the system architecture depicted in Figure 1.1, where two battery banks of different technology (lead-acid as primary and lithium-ion as secondary) and dimensions are the energy storage units. The PV module and the bidirectional Charge Transfer Interconnect bus (CTI-bus) are both managed by a dedicated module, as presented in [4]. Each storage unit is connected to the CTI by means of a bidirectional DCDC converter for level shifting and charge routing, while the PV's one is unidirectional. The CTI itself is connected to the building's supply grid by means of unidirectional AC-DC. The PV module embeds a Maximum Power Point Tracking (MPPT) controller to maximize the energy scavenged. The CTI controller manages the voltages on the CTI-bus and on the bidirectional converters.

On top of the aforementioned HW, runs the Green Energy Controller, an algorithm implemented in Matlab, that computes the optimal control strategy to use with the lead-acid array by using irradiance and load forecasts while the lithium-ion one is used to compensate the error in the forecast in real-time. Since the energy scavenged by the PV module cannot be re-injected into the grid, the optimization problem results in the computation of the optimal charge/discharge sequence of the energy buffers.

2.1 MODELING THE HYBRID ELECTRICAL SYSTEM (HES)

Having a Hybrid ESS [4] allows us to combine the advantages of the two different battery technologies and mitigate their drawback. Lead-acid technology offers better performance in a wider temperature range, has lower price with respect to other technologies and is easier to recycle. Lithium-ion instead performs better in terms of energy density, total number of cycles, charge time and self-discharge rate.

In order to model the behaviour of the HES, a full analytical model has been developed, that describes through equations the behaviour of the Energy Storage System. The development of this model is explained in Deliverable D3.1: Analytical Model of Thermal Behaviour of Data Centre, and the details can be found in the referenced publications [1], [5], [6]. This model allows us to assess the performance of the HES management algorithm.

2.2 THE HES MANAGEMENT ALGORITHM

We have created a discrete-time, two-phase HES management algorithm that firstly computes (off-line scheduling) a long-term plan to use the battery banks, which is then adapted to the real situation (real-time management). The preliminary scheduling scheme, based on Dynamic Programming (DP), manages the lead-acid bank. The real-time controller is an ad-hoc algorithm that compensates the deviation from the forecast by means of the lithium-ion battery bank. The proposed framework implemented in a simulator is depicted in Figure 2.1.

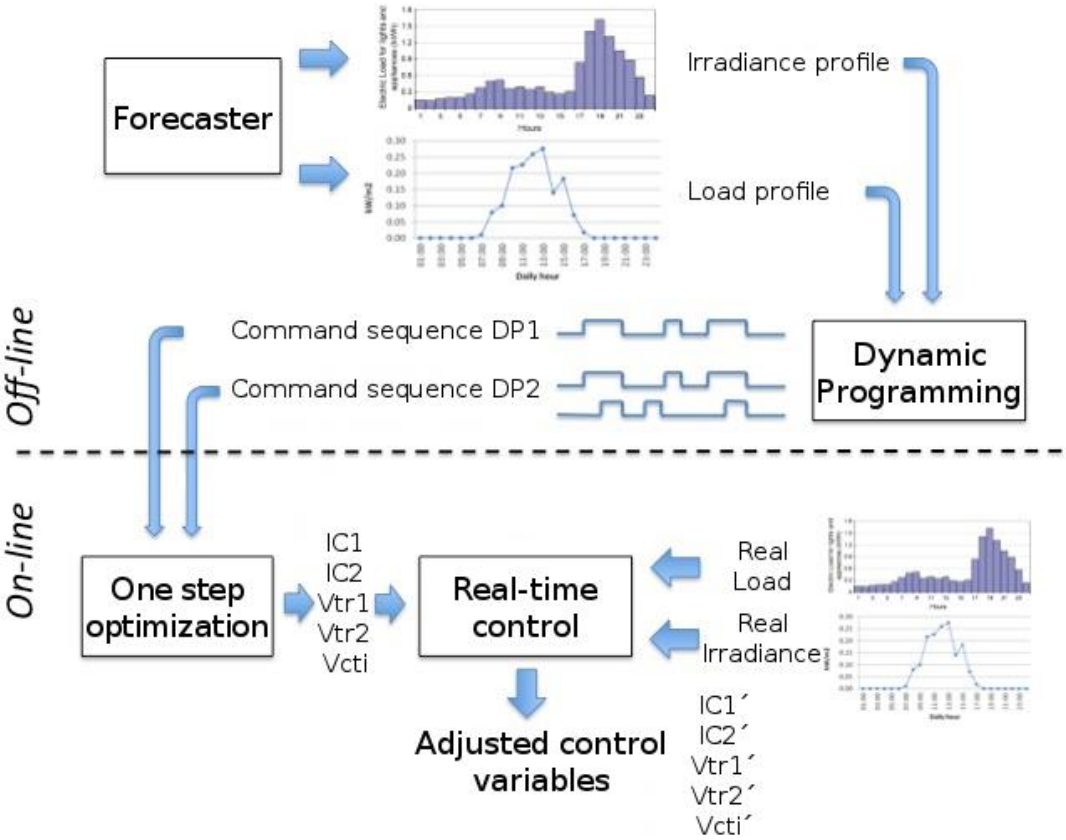

Figure 2.1 – HES simulating framework

We use hourly time-slots to discretize the simulation. During the off-line stage, the simulator computes user load and solar irradiance forecasts to be used as input of the DP. This module then evaluates the optimal charge/discharge profiles for the lead-acid battery bank or both, depending on the simulation goals. The on-line stage optimizes the system control variables by evaluating the difference between forecasted and real power/irradiance profiles. In the same stage the control variables (cf. Figure 1.1 were optimized following the charge allocation scheme presented in [3]. In addition to the two forecasts, the simulator takes as input the horizon of the simulation, the daily price profile for the intake electric energy from the grid and the real data about irradiance and load. In conclusion, we compute the electricity bill amount, the SoH and other battery related parameters to evaluate the performance.

A complete case study, that describes in detail the input/output parameters and evaluates the performance of the green energy controller against two other control strategies typically used in this field, is included in Deliverable 3.2 [1].

3. INFORMATION TECHNOLOGY (IT) INFRASTRUCTURE ENERGY CONTROLLER

The IT Infrastructure Energy Controller is a solution for dynamic power management of modern data centres where scale-out applications represent a significant part of the load. It jointly harnesses server consolidation and v/f scaling in order to reduce the global power consumption while satisfying QoS requirements. The VM allocation algorithm exploits the correlation existing amongst the different applications running in the data centre. Compared to conventional VM placement solutions, the IT Infrastructure Energy Controller provides power savings and improved QoS.

The characteristics of scale-out applications are quite different from traditional HPC workloads in both macroscopic and microscopic scales. At the macroscopic scale, the application, first, is user-interactive; thereby, the amount of required computing capacity is highly variable and fast changing [8] due to the dependence with external factors, e.g., number of clients/queries, etc. Second, the responsiveness (or latency) should come at the first criteria to be satisfied as the level of user satisfaction leads to the success of the business [9]. Finally, the amounts of required CPU and memory resources are usually far beyond the level that a single server can sustain. Hence, massively parallel nodes are cooperatively working by forming a cluster architecture [10]. At the microscopic-scale, among various characteristics of these scale-out applications (see [7]), the memory footprint is far beyond the amount an on-chip cache can sustain; thereby, increasing the on-chip cache size only produces negligible performance improvement.

Because of these aforementioned discrepancies with HPC workloads, existing power management solutions of IT infrastructure of data centres, which neglect or only partially consider the characteristics of scale-out applications, do not exploit all the opportunities to achieve global power savings, as opposed to the IT Infrastructure Energy Controller, specially conceived to be a dynamic power management solution targeting servers that host these new scale-out applications, especially accounting for the correlation information among VMs, while satisfying peak resource requirements. Next, Section 3.1 briefly describes the structure of the correlation-aware placement algorithm, while the specific details are found in Deliverable 3.2 [1].

3.1 CORRELATION - AWARE POWER MANAGEMENT FOR IT INFRASTRUCTURE

The correlation-aware VM placement algorithm is the core of the IT Infrastructure Power Management tool. Roughly, it consists of three parts:

A. The definition of a cost function to efficiently quantify the level of correlation used in the proposed VM placement.
B. The correlation-aware VM allocation scheme while sharing cores among co-located VMs.
C. Provide a way to scale the v/f level to achieve power savings without any QoS degradation.

This problem is a well-known bin-packing problem [11]. To reduce the solution complexity, we apply a First-Fit-Decreasing heuristic. Deliverable 3.2 [1] provides the fine-grained description of the algorithm, along with the models used in the IT Infrastructure Power Management tool to characterize the execution of VMs and the costs to them associated. To further investigate the effectiveness of our proposed controller, a fully detailed experiment is also included, that compares the controller against two other VM-allocation techniques typically used in this field.

4. OVERALL VALIDATION OF GREENDATANET ELECTRICITY CONSUMPTION FORECASTING TOOL

We validated the effectiveness and applicability of our complete ECFT tool (cf. Figure 1.2) to large scale problems using two weeks simulation horizon, workload traces obtained from a real data centre setup and real irradiance profiles. We arranged the simulations in two separate sets, firstly we evaluated the best VMs allocation algorithm in terms of energy performance and QoS; secondly we placed this best scheme into the IT Infrastructure Energy Controller and we executed the joint optimization framework.

4.1 EXPERIMENTAL SETUP

To simulate the IT infrastructure of our urban green data centre workload and energy demand we sampled the CPU utilization of a real data centre setup every 5 min. for one day, then we duplicated the samples up to 14 days. Finally, to generate different samples for each day, we synthesized fine-grained samples per 5 sec. with a lognormal random number generator, whose mean is the same as the collected value for the corresponding 5-minute sample rate.

We modelled the IT infrastructure of a green urban data centre consisting of medium sized facilities with two components: computing power consumption (IT equipment) (i) and computer room air conditioning (CRAC) power consumption as the cooling unit (ii). We guarantee that the power usage effectiveness (PUE) does not exceed 1.3 using modern efficient installations.

Using the utilization traces of a Credit Suisse data centre, we evaluated the effectiveness of the proposed solution with a virtual test bed consisting of 250 servers where the servers are homogeneous. We targeted an Intel Xeon E5410 server configuration which consists of 8 cores and two frequency levels (2.0GHz and 2.3GHz), as in our initial experiments in this deliverable, for consistency reasons.

Similarly, we corrupted a real 14 days solar irradiance profile with random amplitude modulation (1/3 max) and time shifting (2 hours max) to build the irradiance forecast sequence (more details can be found in Deliverable 3.2). At the same time, we used hourly averaged energy consumption profile from the real data centre as forecast, which results in a smoothed profile compared to the original one. This choice is justified from the observation that the data centre workload exhibits an almost fixed consumption pattern independent from weekday and season.

4.2 EXPERIMENTAL RESULTS

In order to select the best VMs allocation scheme for power management to use with the IT Infrastructure Energy Controller we compared three state-of-the-art approaches (see Deliverable 3.2 for details):

- Best-Fit-Decreasing (BFD): a conventional best-fit-decreasing heuristic approach.
- Peak Clustering-based Placement (PCP): a correlation-aware VM allocation which clusters VMs using its Envelope-based correlation classification.
- Correlation-aware VM Placement (Corr): our correlation-aware VM allocation considered as the state-of-the-art approach.

Fig. 4.2 compares the total energy consumption of the three approaches under different number of VMs (obtained by duplicating the trace for 250 VMs) in the system for a horizon of 14 days when we set the V/f level at the time of VM placement, i.e., t_p. The Corr algorithm provides up to 11.6% and 7.3% energy savings compared to BFD and PCP respectively due to using the lower frequency levels more frequently. It is noteworthy that PCP provides almost similar results with BFD because, due to high and fast-changing correlations among VMs in our utilization traces, PCP classifies VMs into only 1 cluster during most of the time periods. When the number of clusters is 1, PCP behaves exactly the same as BFD. Note that the semi-linear trend of the energy consumption depends on the analogous behaviour of the workload among different days, in a typical data centre.

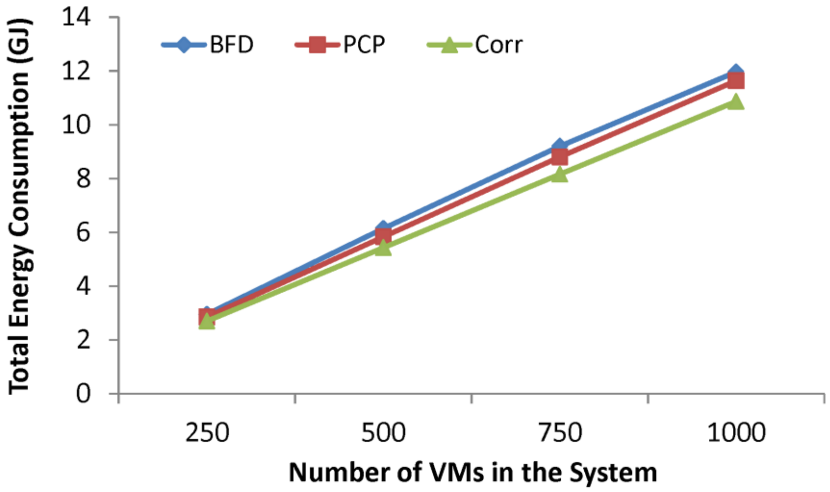

Fig. 4.2 – Total energy consumption of data centre under different number of VMs for a horizon of 14 days

Figure 4.3. shows the maximum violation defined as maximum per-period ratio of the number of over-utilized time instances (i.e., when the aggregated utilization among co-located VMs is beyond the CPU capacity of a corresponding server). As it can be observed, the Corr scheme provides a drastic reduction of the violations, up to 10.4% and 9.6% compared to BFD and PCP respectively. In Corr method, VMs are allocated based on their peak utilizations, which were predicted from their history. Despite the provision based on the peak utilization, we observed quality degradation over the three approaches due to the mis-predictions of the peak utilization, especially during abrupt workload changes under increasing the number of VMs in the system. However, the Corr method can statistically reduce the probability of the violation by co-locating uncorrelated VMs. Thus, the probability of joint under-predictions among the co-located VMs is drastically decreased.

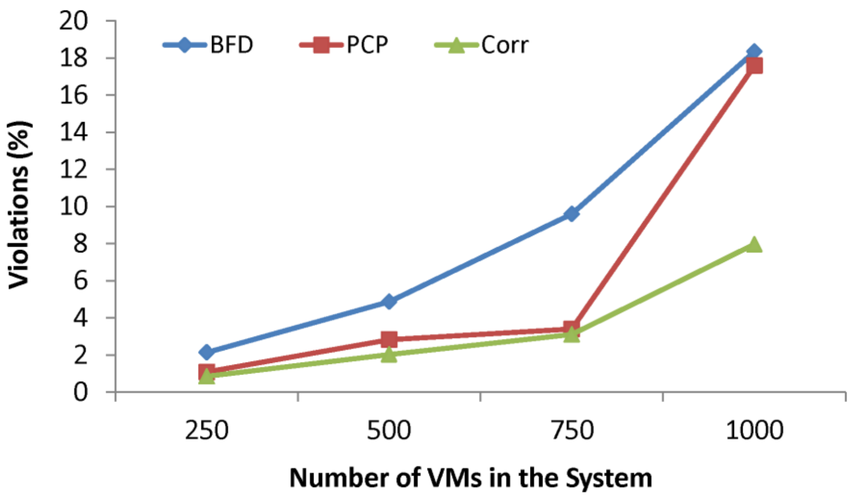

Fig. 4.3 – Trend of maximum violations (%) under different number of VMs for a horizon of 14 days

Using the Corr algorithm, we performed the complete framework simulation (VM allocation, green energy scheduling and communication between the two controllers) with 1 hour time-slots, i.e. t_p = 1 hour, with predictions of upcoming workloads of data centre using a last-value predictor.

Figure 4.4 shows a two days view of the framework evolution with 500VMs and HES-1 configuration. We can observe (Fig. 4.4 **– Two days framework evolution with 500VMs and HES-1 configuration. PV power budget (time-slot average) (top); SoC of the HES (up); workload power supply disaggregation (down); cost per time-slot profile (bottom).**

-top) that most of the green energy provided by the PV module is not used, since the load is not the maximum possible; indeed we can notice (Figure 4.4-middle) that both battery banks can be recharged in few hours. However, this is desirable and beneficial since a typical data centre usually operates in 20–30% utilization rates [10]. By looking at Figure 4.4-up and -down around day 6.8 the role of the constraint on the complete charge/discharge cycles for the primary bank can be observed, this is completely recharged after a partial discharge using the secondary one to minimize the number of cycles and extend its lifetime. Figure 4.4-bottom compares the expenses per time-slot, which result in the end of the 14 days horizon, to an approximate 49% electricity bill saving, more than double that one could get by using PV panels without storage (19%). Similar considerations can be made for the other 3 cases that are not reported for the sake of brevity.

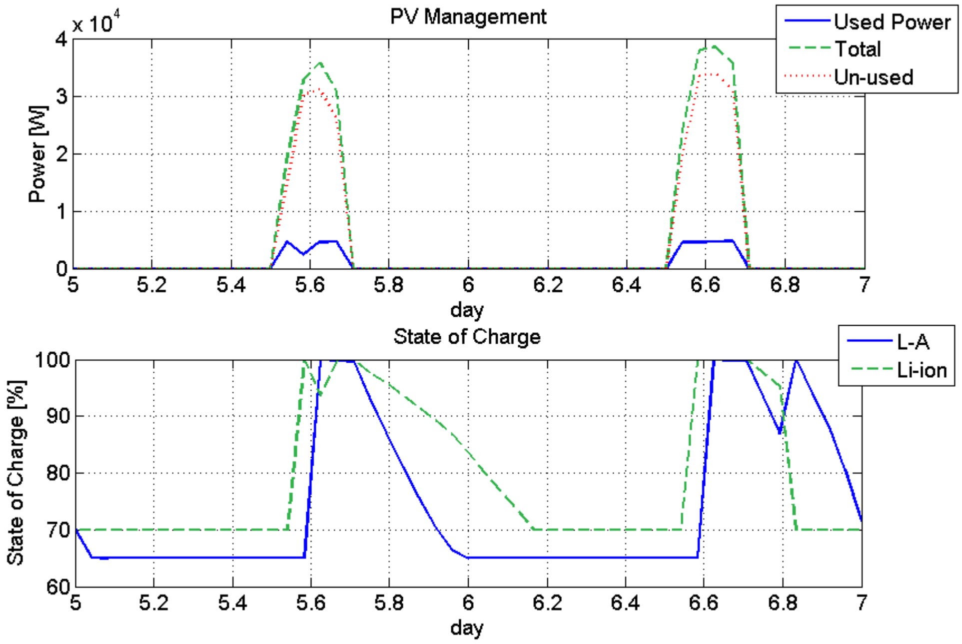

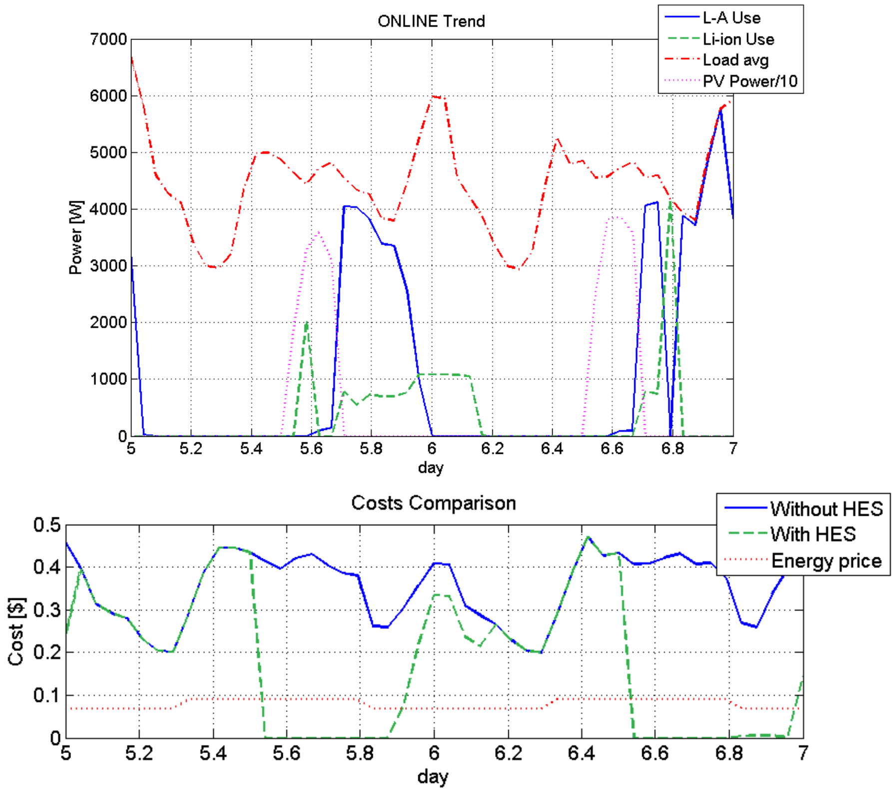

Fig. 4.4 – Two days framework evolution with 500VMs and HES-1 configuration. PV power budget (time-slot average) (top); SoC of the HES (up); workload power supply disaggregation (down); cost per time-slot profile (bottom).

Deliverable 3.2 contains more details regarding the validation of the Electricity Consumption Forecasting Tool.

5. CONCLUSION

This deliverable introduced the GreenDataNet Electricity Consumption Forecasting Tool (further details can be found in Deliverable 3.2 [1]). This tool predicts the energy intake of Green Data centres that run using hybrid electric systems (those containing batteries and photovoltaic modules), and creates an optimal scheduling of tasks in order to minimize the overall energy consumption, while maximizing the use of the batteries (thus minimizing the energy used from the grid).

Testing the ECFT with real inputs from current data centres has shown that the use of this multi-level control scheme developed in GreenDataNet provides significant energy and performance efficiency with respect to state-of-the-art management solutions for the IT infrastructure. Moreover, our projections indicate that the need for greener, more sustainable, IT technologies will decrease the energy consumption of servers and then the presented ECFT tool can provide even more benefits for the future, as new generations of data centres demand more and more energy. As a result, the next step for a holistic management of future green data centres is the integration of this new ECFT tool in the overall Smart Energy Management System (D3.7) that includes the management of the building where the IT data centre infrastructure is located.

6. REFERENCES

[1] **GreenDataNet Deliverable D3.2** - D3.3 - Electricity Consumption Forecasting Tool Design and Implementation

[2] M. Rossi, et al. “Real-time optimization of the battery banks lifetime in hybrid electrical systems”. In Proc. of DATE, European Design and Automation Association, 3001 Leuven, Belgium, Belgium, Article 139 , 6 pages, 2014.

[3] N. Deng, et al., "Concentrating renewable energy in grid-tied datacenters," in Proc. of IEEE International Symposium on Sustainable Systems and Technology (ISSST), 2011.

[4] M. Pedram, et al., “Hybrid electrical energy storage systems,” in Low-Power Electronics and Design (ISLPED), 2010 ACM/IEEE International Symposium on. IEEE, pp. 363–368, 2010.

[5] E. Skoplaki and J. Palyvos, “On the temperature dependence of photovoltaic module electrical performance,” Solar Energy, vol. 83, no. 5, pp. 614 – 624, 2009.

[6] A. Luque and S. Hegedus, Handbook of photovoltaic science and engineering. Wiley. com, 2011.

[7] M. Ferdman, et al., “Clearing the clouds: a study of emerging scale-out workloads on modern hardware,” in Proc. ASPLOS, 2012.

[8] D. Meisner, et al., “Power management of online data-intensive services,” in Proc. ISCA, 2011.

[9] E. Schurman et al., “The user and business impact of server delays, additional bytes, and HTTP chunking in web search,” in Velocity, 2009.

[10] H. Goudarz, et al., “Energy-efficient virtual machine replication and placement in a cloud computing system,” in Proc. Cloud 2012.

[11] J. Kim, et al., “Free cooling-aware dynamic power management for green datacenters,” in Proc. HPCS, 2012.