GAMES: Green Active Management of Energy in IT Service Centres

Massimo Bertoncini[1], Barbara Pernici[2], Ioan Salomie[3], and Stefan Wesner[4]

[1] Engineering Ingegneria Informatica, Italy
[2] Politecnico di Milano, Italy
[3] Technical University of Cluj-Napoca
[4] High Performance Computing Centre Stuttgart

Abstract. The vision of the recently started GAMES European Research project is a new generation of energy efficient IT Service Centres, designed taking into account both the characteristics of the applications running in the centre and context-aware adaptivity features that can be enabled both at the application level and within the IT and utility infrastructure. Adaptivity at the application level is based on the service-oriented paradigm, which allows a dynamic composition and recomposition of services to guarantee Quality of Service levels that have been established with the users. At the infrastructure level, adaptivity is being sought with the capacity of switching on and off dynamically the systems components, based on the state of the service centre. However, these two perspectives are usually considered separately, managing at different levels applications and infrastructure. In addition, while performance and cost are usually the main parameters being considered both during design and at run time, energy efficiency of the service centre is normally not an issue. However, given that the impact of service centres is becoming more and more important in the global energy consumption, and that energy resources, in particular in peak periods, are more and more constrained, an efficient use of energy in service centres has become an important goal. In the GAMES project, energy efficiency improvement goals are tackled based on exploiting adaptivity, on building a knowledge base for evaluating the impact of the applications on the service centre energy consumption, and exploiting the application characteristics for an improved use of resources.

1 Introduction

Over the last years, with the increasing digitalization of the business processes in many application domains, like online banking, e-commerce, digital entertainment, and e-health, the data centre industry has seen a great expansion due to increased need for computing capacity to support business growth. As a consequence, management of IT Processes, Systems and Data Centres has dramatically emerged as one of the most critical environmental challenges to be dealt with and new research directions are being taken towards an energy-efficient management of data centres. An estimation is reported in [15] that the

P. Soffer and E. Proper (Eds.): CAiSE Forum 2010, LNBIP 72, pp. 238–252, 2010.

US servers and data centres consumed about 61 billion kilowatt-hours (kWh) in 2006 (1.5 percent of total U.S. electricity consumption). This estimated level of electricity consumption has been evaluated as equal to the amount of electricity consumed by approximately 5.8 million average U.S. households.

In the last years, large IT systems and data centres are moving towards the adoption of a Service-based Model, in which the available computing resources are shared by several different users or companies. In such systems, the software is accessed as-a-service and computational capacity is provided on demand to many customers who share a pool of IT resources. The Software-As-A-Service model can provide significant economies of scale, affecting to some extent the energy efficiency of data centres. The service-based approach is becoming the most common way to provide services to users, compared to traditional applications developments. Services and their composition, both at the providers' side (to provide new value-added services), and at the users' side (with mash-ups of services composed by the users themselves), are becoming more and more widespread in a variety of application domains. Hence, since the service-oriented approach is steadily increasing for many application domains, its impact on data and service centres will become more and more significant. A very similar model is applied to the provision of services in the High Performance Computing domain where users are allocated to these precious resources in a shared way by using complex scheduling mechanisms.

The EPA report [15] contains a forecast of doubling the energy consumption estimated in 2006 within five years. In this report it is indicated that there is a potential of reducing energy consumption with existing technologies and design strategies by 25 percent or more. However, there is yet a potential for improving energy efficiency in several aspects of a data centre. In fact, despite the big effort that has been put for assessing energy efficiency of IT service centres aiming at the reduction of energy costs [5], the most of these actions have been concerned with solutions in which energy efficiency leverages only on single, yet not interrelated factors, such as the identification of good practices for energy savings based on: (i) the dynamic management of servers according to workload and servers consolidation and virtualization; (ii) the development of low power techniques at IT component level; and (iii) the design of energy-effective facility environments for data centres through reuse of heat or air conditioning. The analysis of the characteristics of the software applications run in data centres are just starting to be considered, such as for instance in the EU best practices for data centres [6].

Mostly, these policies have been implemented in an isolated and fragmented way, not taking into account all the interrelations between the different decision-making layers and were unable to evaluate simultaneous trade-off between power, workload and performance and users' requirements. In particular, the applications running in the service centre are usually only analyzed based on their general characteristics, such as frequency of execution and requests for resources. The analysis of applications at the design level, however, could provide useful information to better manage the resources in the infrastructure. For instance, the structure of the application can be a basis for predicting the resources (e.g., data) that will be necessary for its execution. Such an information can in turn be useful for an internal management of storage resources. On the other hand, information about IT resources can also be used to design energy efficient applications. In fact, while there has been a focus on optimization and negotiation of Quality of Service and performances in the past [11,9], very little attention has been paid to the issues of energy consumption and development of energy efficient services. A first proposal has been presented in [7], where energy consumption and energy efficiency have been considered in composed services at the same level of other quality of service parameters. This allows designing applications that can dynamically adjust to the IT infrastructure state in order to reach energy-efficiency goals, while keeping the agreed quality of service levels.

The vision of the GAMES (Green Active Management of Energy in IT Service centres) project (2010-2012) is for a Green, Real-Time and Energy-aware IT Service Centre. The central innovation sustaining the GAMES vision is that for the first time, to our knowledge, the energy efficiency of the IT Service Centres will be considered simultaneously at different levels, trading-off (1) user and functional requirements and Quality of Services versus energy costs at business/application level, (2) performance, expressed as physical resources workload and Service Level Agreement, against energy costs at IT infrastructure level, (3) HVAC (Heating, Ventilating and Air Conditioning) and lighting versus the power required by the IT infrastructure and the business processes and application, as received by upper levels, at Facility level.

At design time, the assessment and benchmarking of the energy consumption and efficiency of all the different building blocks composing the GAMES IT Service Centres energy efficiency (HVAC, lighting at the facility level, servers, storage, network and processors at IT infrastructure level, services, applications, QoS) will be made for each of the sub-optimal configurations. With this regard, what-if simulation analysis will be carried out in order to determine at design time the best energy-effective distributions of services on the virtualized machines, what will be the best resource and workload configurations with less energy costs, and the impact of these configurations on the energy and carbon emissions balance of the IT Service Centre facility. Historical and required power information and the energy usage profile, combined with Business Intelligence, Data Mining and Information Extraction technologies as well as simulation technologies will be matched with users' business, functional and applications requirements to align energy demand with availability (energy contracted with the utility operator) to design energy efficient applications on an energy efficient infrastructure, able to exploit adaptivity during execution.

The optimized configurations, which will be the output of the GAMES system at design time, will be continuously monitored and adaptively controlled at run-time, through a suitable sensing and monitoring technology infrastructure able to measure temperature, power consumption and humidity of each single IT device (servers, storage, network). The GAMES co-design methodology will aim at co-designing business level applications and services and the IT infrastructure, to support a global energy-aware adaptive approach.

In Section 2 the chosen drivers for the design and the validation of the GAMES approach are described in an abstract way and requirements for the overall approach presented in the following Section 3 are provided. In Sections 4 and 5 we discuss the co-design approach and the adaptive run-time environment respectively.

2 Application Scenarios

The GAMES approach is neither targeting for a specific application sector nor bound to a specific solution for the realisation of the service layer or the applications hosted by the service provider. In order to drive this generic approach with requirements and to allow their validation, two major scenarios are targeted:

- High Performance Computing Service Provision addressing the need for large scale simulations demanding for the co-allocation of a very large number of computing resources for a single task;
- Cloud based computing service provision where the elasticity and dynamism of requested amount of the resources is high.

2.1 Specific Challenges of the HPC Scenario

In this scenario users such as computer engineers submit a job request via a Grid Middleware or directly using an interactive shell environment. Typically such job requests are expressed using the Job Service Description Language (JSDL) [2], an XML based schema allowing to express the requirements such as memory demand, number and speed of CPUs, etc. Such job requests are then either given to a meta-scheduler per provider that aims to find the appropriate resource or directly to a batch oriented queueing system for a single cluster system. In order to allow a differentiation between different consumer types, such requests are typically further detailed with negotiated Service Level Agreements expressing non technical requirements and corresponding guarantees from the provider side as well as penalties in case of violations of the consumer or provider obligations.

From the GAMES perspective, several potentially conflicting constraints need to be actively managed. On one hand, the Service Level Agreements from different customers with different level of importance (in the sense of how important the customer is from a business viewpoint) and different penalties in case of violations need to be aligned with policies like "prioritize large jobs compared to small ones for this resource", "maximise utilization of the resource" with Green Performance Indicators such as "reduce workload during days where free cooling is impossible due to high outside temperatures".

As a further illustration consider the simplified example of a computing resource having 20 nodes. Currently 3 jobs requiring 15 computing nodes each and a couple of smaller jobs with only 8 nodes each are waiting in the job queue.

The policy of preferring large jobs would prioritize the two 15 node jobs ahead of the 8 node jobs. The policy for reducing the workload if the outside temperature does not allow free cooling (and therefore higher costs) would switch off 12 nodes during the hot time of the day pushing the 8 node jobs and would move to the 15 node jobs during the evening/night when free cooling is possible again. The job mix is typically much more varying and the number of nodes on high end production systems is several thousands.

An additional problem related to this scenario is that many applications hosted by the provider or provided by the consumer at job submission time could be even only available in binary form (e.g. provided by an Independent Software Vendor) and cannot be easily made GAMES enabled to react on commands that are control attempts from the GAMES framework. Consequently, the GAMES framework can only intervene with the supporting infrastructure such as the queuing systems and job scheduling frameworks and the SLA assessment and validation infrastructure.

Nevertheless, by influencing on one hand the underlying provisioning middleware as well as influencing the queuing infrastructure, a plethora of control possibilities do exist. For example one can limit access to certain queues allowing prioritised access to the resources for certain users or jobs, as well as limit the maximum amount of time one can use a resource demanding a restart capability of the application developer to continue a simulation after it got canceled at the last stable state. Additionally, different customer profiles provide additional control possibilities. Such customer profiles could range from cost optimized best-effort computing, over time boxed simulations (e.g. in a car design process) up to urgent computing cases where simulation results are expected to be used as input to a decision support system for a medical treatment.

2.2 Specific Challenges of the Cloud Scenario

In contrast to the scenario above, the applications hosted in this scenario have no pre-defined maximum wall clock time or a fixed amount of resources but demand a dynamically changing amount of resources (commonly referred to as elasticity) from the provider and do not have the requirement to use several thousands of computing nodes exclusively at the same time but can operate on top on a virtualized infrastructure potentially sharing one physical computing node with several other customers running their own virtual machines. Additionally, properties like reliability and robustness are of much higher importance as in contrast to the simulation jobs of the HPC scenario it is not easily possible to perform a re-start of the simulation at the last checkpoint some hours later without affecting the Quality of Experience (QoE) of the consumer.

Considering that beyond the Infrastructure as a Service (IaaS) paradigm, also Platform as a Service (PaaS) and Software as Service (SaaS) are becoming more and more important in the cloud provisioning model (see also [10]), one can also assume that in this case the applications that are supposed to be provided as services can be fully GAMES enabled and provide a direct interface from the cloud operation level over the platform up to the software level.

However, similar to the HPC use case, the differentiation of the Service Level for different customers, business policies to prefer specific kinds of applications (e.g., applications with a high demand or low demand for elasticity) are similarly conflicting with the demand to save energy.

2.3 Major Requirements Summary

Summarising the major challenges of the two distinct scenarios above, one can say that the following aspects are currently missing in existing data and computing centre infrastructures:

- A mitigation framework allowing to derive clear policies and actions for the underlying infrastructure based on the potentially contradictory and conflicting optimisation goals of the different aspects (e.g., SLAs, Green IT aspects, Differentiated customer support, and so on).
- A rich sensing and monitoring framework allowing the collection of the necessary information on all levels from the facility and environment, over physical resources over platform services up to the applications themselves.
- Data Mining and Reasoning elements analysing the historical data collected aiming for a prediction and derived counter measures to bring the overall infrastructure back to the desired operational point.
- Actors and control elements on the different levels allowing the counter measures to be applied on the appropriate level in the necessary short time.

Orthogonal to this more operational oriented requirements, one can clearly realise the need to plan for the monitoring at the design phase of a building, hardware and the corresponding software services. Similarly, one needs to think about which elements can be controlled and which hooks need to be designed into this overall setting. Beside a Runtime Environment aiming for the optimisation of the operation of the application services, similarly the right design of the overall system from facility and energy provision model up to the service composition and application layer is of equal importance.

3 The GAMES Approach

In the GAMES approach, we consider a joint management of the applications and the IT infrastructure on which they are running. Both at the application level and at the infrastructure level, we assume that the system is adaptive: the applications can change their modes of operation at run time, and the IT infrastructure can be dynamically reconfigured in the service centre. Adaptation is performed according to a number of adaptation strategies that are encapsulated in adaptation rules that are evaluated at run time. The adaptation strategies and their rules are designed for the service centre taking into consideration both the IT infrastructure and its capability of providing an autonomic behavior and the characteristics of the applications being executed on the system and their requirements. We assume that the service centre is exploiting a virtualization of the IT infrastructure, so that virtualized IT resources can be associated to each application, thus application management and IT management can be decoupled. The adaptivity of the system is performed considering its general context of execution, which includes physical environment parameters such as the external temperature and humidity, the parameters of the service centre facility infrastructure, which include cooling and heating and servers and storage parameters, and application parameters, mainly in terms of Quality of Service requirements and utilization rates of infrastructure components assigned to the application. We propose an adaptive SBA (Service Based Architecture) as the basis of the energy-aware design and management of service-based information systems and their IT infrastructure. At the application level, we consider applications as being executed by invoking services or composite services, which can be possibly dynamically modified. To each service, Service Level Agreements are associated, covering Quality of Service requirements, which can be dynamically renegotiated. Similarly, platform services and infrastructure are all considered as services in the system, creating a complex environment of service compositions, where each service is associated to a number of possible adaptation actions.

The goal of the GAMES approach is to realise a self-adaptive data and service centre architecture across all kinds of offered resources ranging from IT infrastructure data and facility over computing up to the service layers. The conceptual architecture in Figure 1 shows the components needed for run-time management to continously balance the agreed service contracts and to derive the necessary measures needed based on the monitored values (energy consumption, load situation, risk to fail on an SLA, etc.) as well as the interaction with the design environment.

In the following, we briefly introduce the main components of the GAMES architecture:

– the Energy Sensing and Monitoring Infrastructure (ESMI)
– the Run-Time Environment (RTE)
– the Design-Time Environment (DTE)

The **Energy Sensing and Monitoring Infrastructure** (ESMI) provides services to interact with the energy grid, with the environment monitoring infrastructure and with the data centre resources, for energy consumption and other physical measures. The ESMI has an energy service layer providing basic monitoring, messaging, event derivation features, and mining services for analysing historical data targeting the generation of useful adaptation patterns and knowledge. The ESMI will be partially based on the energy service layer being developed in BeAware [8]. The sensing infrastructure will be interfaced with monitoring services, which will in addition gather relevant information from the IT infrastructure and SBA layer, generating relevant events from the sensor information. A context management support module will manage context information.

The **Run-Time Environment** (RTE) provides an energy-aware and self-* adaptivity controller. It includes functionalities for event analysis, based on

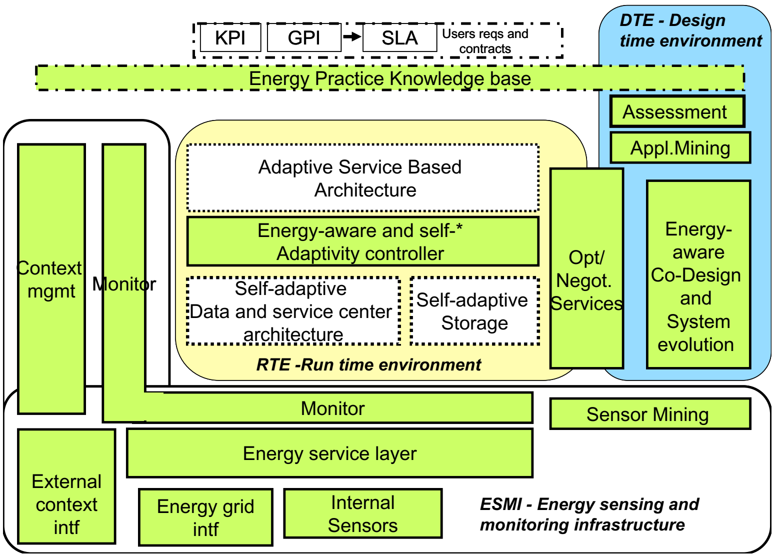

Fig. 1. GAMES architecture

the general knowledge of the environment and energy characteristics of services, controlling the adaptivity under a global perspective of a service and an architectural level. Control will be also based on a general optimiser and negotiator, which, starting from static tools for architecture optimisation and SLA templates, will be enhanced with dynamic and energy-aware functionalities, exploiting also the Energy Practice Knowledge Base. The self-adaptive data centre architecture module comprises an adaptation of the architectural part and of the storage-part through strategies and decisions on data placement and storage quality of service based on access patterns and mapping of application services to data storage level.

The **Design Time Environment** (DTE) will support an energy-aware co-design of service-based information systems and IT architecture in the data and service centre. All design choices are driven by users demands expressed as a set of Key Performance Indicators (KPI) and Green Performance Indicators (GPI) that are part of the negotiated Service Level Agreements (SLA). Starting from a static evaluation of existing configurations, optimisation and negotiation techniques for design time, choices will be developed to devise the optimal functioning points to be exploited for run-time adaptivity. The design will include also the identification of the observable needs for optimal and efficient run-time event detection. Users involvement will be considered through test cases and user experience models. An assessment tool will provide an initial analysis of the users requirements, service and data characteristics and IT infrastructure and facility from which the energy-aware adaptive service and infrastructure design will start.

4 Designing an Energy-Efficient Service Centre

Energy-aware service-based information systems design will be tackled based on a three-fold perspective: a) strategic-level decisions in developing green IT service centres (e.g., identifying Green Key Performance Indicators (GPI) and analysing the impact of QoS business process levels on energy costs); b) developing control strategies to evaluate, optimise, and control services and data at run-time on multiple time scales and adapt them at run time; c) realizing technological mechanisms and tools to reduce the energy consumption of IT service centres based on self-adaptive services and architectures. Energy savings can be obtained by exploiting the characteristics of existing adaptive platforms both at the business/application level, where adaptive service compositions can be executed, and at the architectural level, based on adaptation of IT architectures and components. The problem to be solved is how to combine the existing approaches in a layered architecture, considering a large number of information systems using the same services and sharing the same data centre(s). We propose a combined design-time and run-time approach. At design time, co-design is proposed to create adaptive service-based information systems and self-adapting architectures based on the requirements. At run-time, we propose an event-based adaptation process that takes into consideration the run-time context information (energy consumption) and design-time context information (user and business contexts).

We will focus on the design of energy-aware information systems, in which the information system functionalities and the IT system architecture are co-designed to get improved energy efficiency. The energy dimension is currently not considered in information systems design, where functionality and quality of service considerations are driving design choices. Based on some research experiments and simulation [1,3], we advocate that considering the energy consumption dimension, different and more efficient design choices could be performed.

Examples of energy-aware co-design include several aspects at different levels: strategies to minimize the number or similar/redundant services, e.g. by using virtualisation technologies or a balanced number of servers performing supporting services operations (e.g., having only a minimal number of authentication servers); an evaluation of the impact on needed cooling capacities based on different load scenarios of servers; a focus on business process analysis of core activities-services-data as shown in [14], annotating business processes with meta-information useful for performing an energy-efficient management of the application at run time.

We will develop a cost-based approach to design the system globally and to select the adaptation strategies that are recommended at run time at the application (process/service composition) and at the IT level and to identify the variables and components which need to be monitored in order to ensure a correct control of the system. Business processes will be analyzed considering their quality of service requirements and their needs for IT processing infrastructure. Process meta-information will include data requirements and task dependencies, the ability to use alternative services in service compositions, and their context-awareness, in order to be able to enhance the adaptive capability of the application itself, both at application management and IT management level.

5 Energy Efficiency at Run Time

GAMES defines a new approach for a run-time, energy-aware adaptive mechanism. The basic concept is to consider and use the system context situation enhanced with energy/performance information for controlling / adjusting / enforcing the run-time energy efficiency goals. We approach the problem of minimizing the energy consumption in a service centre by using Dynamic Power Management (DPM) and consolidation techniques. In a classical data centre, the computing resources are over provisioned to handle the peak value. We propose to minimize the energy consumption by monitoring the service centre servers to determine the over provisioned resources with the goal of putting them in low power states.

Layered feedback architecture will be considered for run-time controlling of systems performance/energy ratio, by combining autonomic and context aware computing methodologies, techniques, algorithms and tools with methods and tools specific to the systems and control theory.

We propose the development of two types of control loops that will be used to adjust and adapt the system execution to the energy efficiency goals established in the co-design phase: a set of local control loops associated to IT Infrastructure servers and one global control loop associated to the whole system.

5.1 Local Control Loop (Server Level Controller)

The local control loops are used to locally optimize the energy consumption at the level of each server. The controller is developed by using a set of server specific energy optimization rules, predefined at design time, which can be executed on a very fine time grain, without affecting the system overall performance. The main guiding idea of server level energy efficiency is to pro-actively identify and take appropriate actions to reduce servers resource over-provisioning so that it matches the application requirements. This way, energy can be saved by transitioning the over-provisioned resources to low power states while maintaining performance by satisfying the application requirements. For the local loop controller design we have used server specific DPM and workload allocation techniques, considering the CPU and external storage as the main controlled resources. The local loop controller takes DPM actions based on the current server state and on the workload received from the global control loop (see Fig. 2). The local loop controller uses the context model instance, a knowledge-base and a prediction engine for inferring the most appropriate DPM actions.

CPU management. For the energy-aware CPU management, two workload allocation strategies and an improved Dynamic Voltage Scaling (DVS) technique is proposed. The first workload allocation technique distribute as many tasks per core (allowing the unused cores to be switched into low power states) and per

CPU (allowing the unused CPUs to stay in idle states as much as possible) while the second strategy distribute evenly the tasks and obtaining as a result a lower overall utilization, thus enabling the transition of the entire core / socket system to lower frequencies. Starting from the observation that service requests are unpredictable over time and peak loads are orders of magnitude larger than those in steady states, we have decided to use a fuzzy-logic based control algorithm, by adapting the work presented in [4] for implementing the DVS techniques aiming at processor dynamic voltage scaling.

*S*erver storage management. The local control loop also manages the server HDDs aiming at maximizing energy efficiency. Unlike the CPU which is very flexible in terms of power management, the state transitions of HDDs are more rigid, involving significant performance degradation when executing a wrong power management decision. For HDD management, the local control loop uses an advanced adaptive algorithm, based on adaptive learning trees [12], aiming at identifying the access patterns and deciding about the possible spin-down of the drives. Moreover, the resource management component of the local control loop will attempt to distribute the incoming workload associated storage requirements in such a way to maximize the idleness periods of some drives while preserving the SLA indicators.

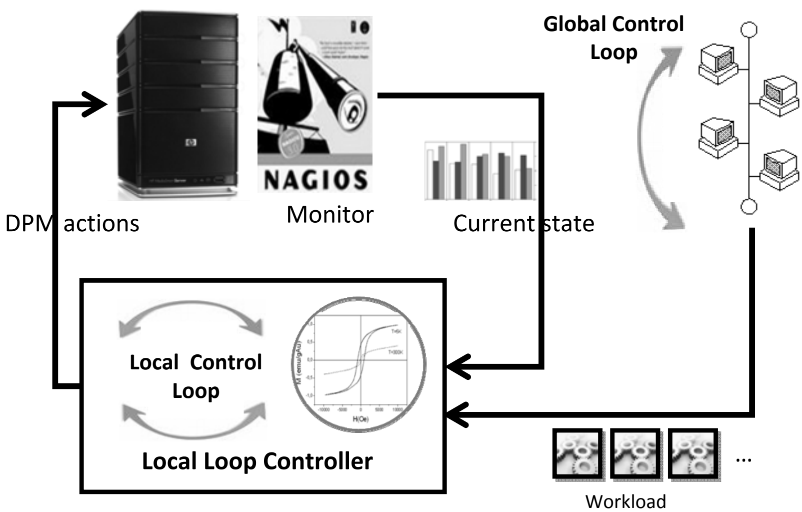

Fig. 2. Local control loop

Local control loop architecture for a virtualized server is presented in Fig. 3. Virtualization allows for better control of energy savings by (i) the fine grain resource allocation in a per virtual machine policy and by (ii) facilitating server consolidation by virtual machine migration, leading to the complete shutdown of unused servers. Local control loop consists of three main modules: the resource allocator, the monitoring subsystem and the Dynamic Power Management (DPM) module. The resource allocator module assigns resources to the virtual machines so that it delivers the expected performances while the CPU and HDDs make low power states transitions as often as possible. The monitoring and analyzing module analyze the state of the server and the state of the virtual machines and provides the monitoring data to the resource allocator module and to the DPM Controller. The DPM module controller enforces the power states transitions determined by the CPU and HDD prediction engines.

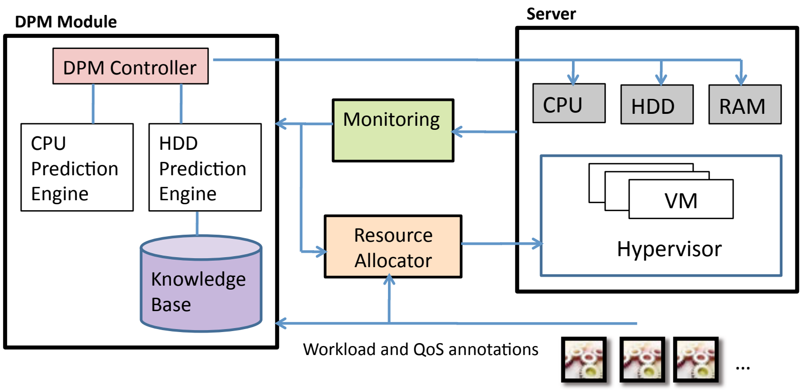

Fig. 3. Local loop controller architecture

5.2 Global Control Loop (Service Centre Level Controller)

The objective of the global control loop (see Fig. 4) is to evaluate the energy enhanced service center context situation (represented by the context model instance) and determine the most appropriate adaptation decisions to enforce and realize the Key Performance Indicators (KPI) and GPI (Green Performance Indicators) defined in the co-design phase. The context model instance records current workload and energy data from all service centre individual servers, environmental data and business process features. The context model used in the GAMES project is based on the RAP model [13] which represents the context as three sets together with their relationships: a set of resources, a set of actions and a set of policies. Context resources are context data generators or consumers and can be classified as passive or active context resources. Passive context resources (such as physical or virtual sensors) aim at capturing and storing context data while active context resources (such as actuators or facilities) can interact directly with the context and modify the context state. Context policies are used to represent the set of conditions used to guide and control all interactions within the context. The set of adaptation actions are specified in design time and are executed in run-time to enforce the context policies.

The global control loop objectives are implemented as an energy-aware and self-* controller by using a MAPE cycle paradigm which involves the Monitoring, Analyzing, Planning and Execution phases. During the monitoring phase, the context instance is updated with workload and energy related data collected from service centre computing resources and facility, as well as with QoS and SLA related data of the incoming business process requiring execution. In the analysis phase the context instance is evaluated to determine if the predefined context policies are fulfilled. If there are no broken policies, the system is considered as being in an energy-efficient state and no adaptation action is necessary. If one or more policies are broken, an adaptation plan consisting of a set of healing actions is generated in the planning phase. The execution phase is running the adaptation plan actions, aiming at system transition into an energy-efficient state.

At the global control loop level we have defined three types of adaptation actions: computing resource oriented adaptation actions, facility resource oriented adaptation actions, and application oriented adaptation actions. Computing resource oriented adaptation actions are executed to enforce the set of predefined GPI and KPI indicators on the service centre computing resources. We defined two types of such actions: consolidation actions and DPM actions. Consolidation actions aim at identifying computing resources featuring inefficient workload distribution and taking balancing actions to distribute it in an energy efficient manner. Consolidation actions are implemented as activity migration or activity deployment actions. DPM actions aim at determining the over provisioned resources with the goal of putting them into low power states for minimizing energy consumption. Facility resource oriented adaptation actions such as adjust the room temperature or start the CRAC, can be enforced through the service centre facility active resources. Application oriented adaptation actions are executed upon the GAMES-enabled applications targeting their redesign for energy efficiency.

The global control loop decisions may also include other energy-aware context-based adaptivity actions such as minimizing the necessity of calling a remote service when one local similar service is available (minimize data/service transfer), minimizing the substitution of services during maintenance, optimizing the number of necessary backup operations or favoring the use of services that require low energy.

To determine the best adaptation actions for a given context situation, the service centre level controller uses reasoning / learning and data mining tools, what-if analysis tools and a knowledge base. To derive knowledge about the service centre and its energy efficiency, the GAMES framework integrates information models that uniformly represent the system historical energy consumption related data. The general approach is based on extracting domain knowledge base from large amounts of historical data by using data mining techniques. The historical energy consumption related data will be also used together with a traceability model to understand the impact of changes in the provisioning infrastructure on energy efficiency and service quality, in order to allow both operators and consumers to select the appropriate mix as needed. With the GAMES framework it will be possible to align business requirements such as optimize for low power demand providing response time up to 200ms versus optimize response time based on historical data and the currently monitored status. By combining at design and run time the historical, predictive, context

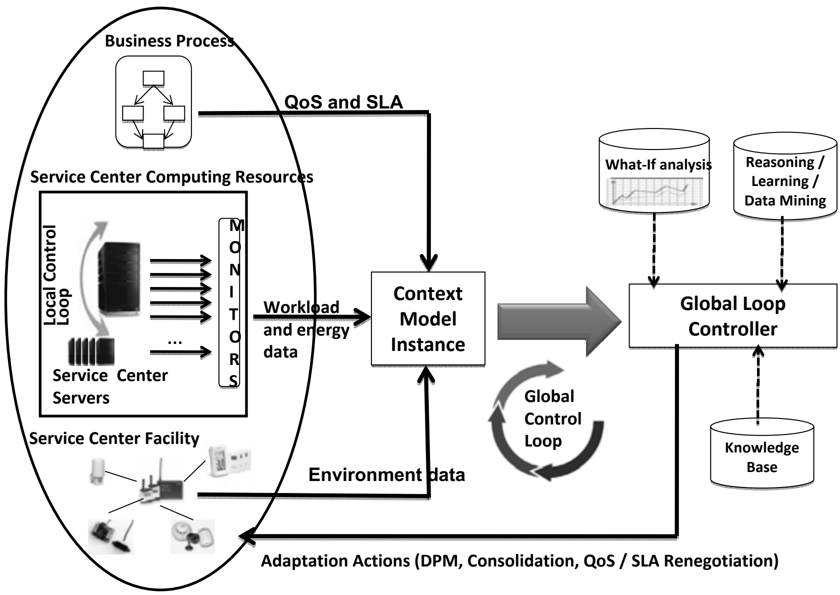

Fig. 4. Global control loop

and the externally available information with the GAMES knowledge base will allow the selection of the most adequate adaptation patterns and profiles.

6 Conclusions

This paper has presented the GAMES approach to design and manage energy-efficient service centres. For implementing in a successful way the GAMES concept of energy efficiency, new overall energy efficiency metrics are needed, which will be able to assess the energy efficiency and carbon emissions in an integrated way, combining the facility with the business/process and IT architecture levels, while the most popular ones nowadays (PUE and DCiE, defined by the GreenGrid consortium [5]), are dealing only with the facility level. With this regard, the GAMES project will define and introduce new energy efficiency and emissions metrics, the GAMES Green Performance Indicators.

The general approach of co-design and adaptivity both at service and at infrastructure layer need validation, both from a theoretical point of view and from experimentation. Models and tools to be developed must be sufficiently performant and the monitoring light enough not to overload the running system. Validation in the project is planned within two large data centres, on experimental settings.

Acknowledgments. This work has been partially supported by the GAMES project (http://www.green-datacenters.eu/) and has been partly funded by the

European Commission's IST activity of the 7th Framework Program under contract number ICT-248514. This work expresses the opinions of the authors and not necessarily those of the European Commission. The European Commission is not liable for any use that may be made of the information contained in this work.

References

1. Almeida, J., Almeida, V., Ardagna, D., Francalanci, C., Trubian, M.: Managing energy and server resources in hosting centers. In: Proc. ICAC (2006)
2. Anjomshoaa, A., Brisard, F., Drescher, M., Fellows, D., Ly, A., McGough, S., Pulsipher, D., Savva, A.: TheJob Submission Description Language (JSDL) Specification, Version 1.0. Technical report, Open Grid Forum (2005)
3. Ardagna, D., Cappiello, C., Lovera, M., Pernici, B., Tanelli, M.: Active energy-aware management of business-process based applications. In: Mähönen, P., Pohl, K., Priol, T. (eds.) ServiceWave 2008. LNCS, vol. 5377, pp. 183–195. Springer, Heidelberg (2008)
4. Chen, Z., Zhu, Y., Yuan, M.Y.: Towards self-optimization in utility computing using fuzzy logic controller (2005)
5. Belady, C. (ed.): Green grid data center power efficiency metrics: PUE and DCiE (2008)
6. EU Stand-by Initiative. 2010 Best Practices for the EU Code of Conduct on Data Centres. Technical report (December 2009)
7. Ferreira, A.M., Kritikos, K., Pernici, B.: Energy-aware design of service-based applications. In: Baresi, L., Chi, C.-H., Suzuki, J. (eds.) ICSOC-ServiceWave 2009. LNCS, vol. 5900, pp. 99–114. Springer, Heidelberg (2009)
8. Gamberini, L., Jacucci, L.G., Spagnolli, A., Bjorkskog, C., Kerrigan, D., Chalambalakis, A., Zamboni, L., Valentina, G., Corradi, N., Zappaterra, P., Bosetti, G.: Technologies to improve energy conservation in households: The users' perspective, Maastricht. In: First European Conf. Energy Efficiency and Behaviour (October 2009)
9. Hasselmeyer, P., Koller, B., Schubert, L., Wieder, P.: Towards SLA-Supported Resource Management. In: Gerndt, M., Kranzlmüller, D. (eds.) HPCC 2006. LNCS, vol. 4208, pp. 743–752. Springer, Heidelberg (2006)
10. Jefferey, K., Neidecker-Lutz, B.: The Future of the Cloud (January 2010)
11. Koller, B., Schubert, L.: Towards autonomous SLA management using a proxy-like approach. International Journal of Multiagent and Grid Systems 3, 313–325 (2007)
12. Lu, Y., Chung, E., Simunic, T., Benini, L., Micheli, G.D.: Quantitative comparison of power management algorithms. In: Proceedings of the Conference on Design, Automation and Test in Europe (2000)
13. Salomie, I., Cioara, T., Anghel, I., Dinsoreanu, M.: RAP - A basic context awareness model. In: Proceedings of 4th IEEE International Conference on Intelligent Computer Communication and Processing, Cluj-Napoca, Romania, pp. 315–318 (2008)
14. Schmidt, N.-H., Erek, K., Kolbe, L.M., Zarnekow, R.: Towards a Procedural Model for Sustainable Information Systems Management. In: Proceedings of the 42nd Hawaii International Conference on System Sciences, HICSS 2009, Hawaii, USA, pp. 1–10. IEEE Computer Society, Los Alamitos (2009)
15. U.S. Environmental Protection Agency. Report to congress on server and data center energy efficiency public law 109-431. Technical report, ENERGY STAR Program (August 2007)