Business Process Design: Towards Service-Based Green Information Systems

Barbara Pernici, Danilo Ardagna, Cinzia Cappiello

Politecnico di Milano, piazza Leonardo da Vinci 32, Milano, Italy

Contact author: barbara.pernici@polimi.it

Abstract: This paper discusses the impact of energy consumption on information systems and business processes design. The goal is the development of context-aware and sustainable information systems where energy consumption reduction is considered at the technological level, on the basis of adaptable technology, on the governance level, with the design of context-aware processes and data, and at the strategic level. The adoption of a service-oriented approach and the interconnection among the different levels are discussed.

1. Introduction

Climate debate and Sustainable Growth concern over energy use have thrown the spotlight on power consumed by Information Technology (IT). Since information management has become pervasive in modern society, the impact of IT on energy budget is becoming more and more significant. The growth in the number of servers and the increasing complexity of the network infrastructure has caused an enormous spike in electricity usage. Nowadays, service centers alone consume 1.5% of the power produced in the US, and are projected to reach 4.5% within 5 years. This has driven efforts to reduce the power consumed by service centers and at the device level. Another issue related to energy consumption which is influencing IT decisions in many countries has been the introduction of carbon taxes connected with CO_2 emissions and the carbon credits debate. There is a strong need for new instruments able to analyze the impact of IT decisions and enterprise business processes on energy and costs and to set an active control on Information Systems (IS) on the basis of the energy driver.

In a recent paper, (Murugesan 2008) analyzes the environmental impact of IT along four dimensions: design, use, manufacturing, and disposal of IT systems. This paper analyzes mainly the first two dimensions: design of Information Systems under an energy consumption perspective, focusing on service and information management, and use of Information Systems, focusing mainly on the reduction of the resources needed for processing information and for information storage after its elaboration.

Several research directions are emerging in the literature to improve the quality of Information Systems. Software as a service suggests economies of scale and, in recent years, large service centers have been setup to provide computational capacity on demand to many customers who share a pool of IT resources. With the development of the Service Oriented Architecture (SOA), multiple service providers can offer functionally equivalent web services which may differ for their offered Quality of Service (QoS). Virtualization technology can be exploited to improve the efficiency of use of hardware infrastructures, increasing the average percentage utilization of the systems and allowing the dynamic reconfiguration of physical infrastructures.

Nowadays, IT systems are required to provide high availability, reliability, and performance. Increased business continuity requirements, which necessitate the replication of data to protect against system failures, and the growing trend to leverage a great variety of information for business operations, are driving customers to store more data than in the past. IDC predicts storage annual growth around 60%. Redundancy improves systems' QoS, but may introduce energy inefficiencies.

In our work, we focus on the point that energy efficiency should be given a very relevant role in Information Systems design. In fact, advances in autonomic and self-healing service-based systems enable a potential reduction of system redundancy, and energy optimization related to data management is more and more challenging.

The goal of the paper is to discuss a systematic set of theoretical approaches and methods and tools towards the implementation of "energy-proportional" Information Systems, i.e. to consume energy according to the amount of work performed, contrary to the current state of the art, in which systems are mainly under-utilized and perform many non-essential tasks. Our focus is mainly towards a green approach to the use of information and application in Information Systems, at a software and management level, rather that on the development on the underlying hardware infrastructure. The aim is to reduce redundant components by designing systems able to react dynamically to variable operating conditions and maximize the use of resources with minimal energy consumption, through their self-reconfiguration. Current work in sustainable and power aware computing suggests designing systems with a trade off between QoS and energy consumption.

An information system approach allows considering all components which are needed for information management in an organization in a global framework. We will aim to develop an "IS purifier" approach, in analogy to what is being done for cleaning water for a sustainable environment. The approach will be based on a triple perspective approach to reduce and streamlining energy consumption in Information Systems: a) IS green governance to support strategic decisions in developing green Information Systems and to evaluate the impact on energy consumption and CO_2 emissions, b) IS green control to evaluate, optimize, and govern services and data both at an application and at an infrastructural level with

emphasis on energy efficiency, and c) specific techniques to support   the purification of parts of the system (e.g., service and data redundancy elimination) when critical situations arise that might have a deep impact on energy consumption.

2. Related work

Green computing is a new discipline and practice of using computing resources with a focus on the impact of IT on the environment. Initiatives have been taken in order to reduce the use of hazardous materials, maximize energy efficiency, and recycling products and materials. Green computing has been started in 1992 with the foundation of the Energy Star program and currently is sustained both by government regulations (e.g., EU directives 2002/95/EC on the reduction on hazardous substances, 2002/96/EC on waste of electronic equipment) and by the industry (e.g. EPEAT criteria for the evaluation of products by the Green Electronic Council, or the IBM Big Green project on the energy efficiency of data centers).

   At the governance layer, literature contributions are by (Jokinen 1998), which examines the relationship between the information society and sustainable development on theoretical and conceptual levels, and (Cohen 1998) which focuses on e-commerce and environmental sustainability. Though no definitive conclusions are drawn, these papers provide a starting point and raise green IT issues at an early stage. (Yi 2007) provides a review of the state of the art of projects, literature, and products addressing how e-business/ICT affects the environment. They mention ISA, an assessment tool to examine approaches and activities to evaluate the implications of information society technologies (in the EU 5th Framework Pro-gram), and a few scenarios, criteria and indicators as tools for identifying environmental impacts inherent in the information society. (Williams 2008) advocates the need for creating an ecosystem map, introducing the concept, to be developed, of Key green performance indicators. (Borup 2007) discusses the increasing focus on foresight exercises as a tool in public governance of research and industrial innovation systems. Technology foresight, as a systematic analysis and discussion about possible technology futures, gives evidence to Green technology foresight (GTF), or environmentally oriented technology foresight. Environment issues and thoughts about sustainability are if not high on, then at least visible, on the agenda in society in general. This has created a new area of relevance and legitimacy for techno-scientific development and research and new opportunities and markets for environment technology are gaining importance. An integration of science and technology policy with environment policy seems more relevant and obvious than earlier. These things are reflected in the appearance of green technology foresight and also in the inclusion of environment aspects in more general technology foresight activities.

A prospective simulation study of a project commissioned by the EU Institute for Prospective Technological Studies models all known relevant ICT effects using a system dynamics approach in combination with scenario techniques and expert consultations. The results offer debate on how the chosen environmental indicators will have very variable absolute values in 2020, which will depend on many uncertain model parameters. This paper contributes to the general understanding of the environmental impacts of ICT and provides a useful basis for policy-making in the fields of ICT and environment. Scenarios are still qualitative planning and communication tools, but a need arises for predictive management tools. The life cycle assessment (LCA) tool, known as life-cycle analysis, allows for estimating the total environmental impact of a given product or service throughout its lifespan. LCA, as an extended economic or environmental input–output life cycle assessment, has been the most common method applied so far, since awareness and understanding of governance issues are still rather un-studied. An artificial neural network approach embedded in a predictive and empirical model is reviewed as a tool for companies to analyze how their e-business operations influence the environment, based on indicators to be chosen to represent e-business/ICT factors and geo-environmental performance.

At the technological layer, energy consumption in computing devices has received great attention by the research community. Much work has been done in order to achieve power reduction in mobile systems (both at device and network level (Kremer 2001)) to extend battery life, while low power techniques are now progressively introduced in server environments. All major industrial players are taking a position in the green IT area (see for instance (Williams 2008) from Microsoft, (Barroso 2007) from Google, and the IBM Big Green Project launched in May 2007. In academic research, to reduce energy consumption, autonomic techniques have been developed to manage workload fluctuations (Chase 2001) and to determine optimum trade-offs between performance and energy costs by switching servers on and off and by implementing DVS mechanisms (i.e., reduction of CPU frequency of operation of new servers) (Ardagna 2007c, Quin 2007, Kusic 2006). Overall, the development of energy-proportional computing systems is advocated (Barroso 2007), since in current implementations the energy consumption vs. performed work dependency is highly non-linear.

In the Software as a Service area, research has been performed on the maximization of the QoS for end users. First solutions proposed greedy approaches (Maamar 2003) which select one at a time, the best candidate service suitable for the execution of business process activities. (Zeng 2004, Ardagna 2007b, Canfora 2005, Yu 2007) propose more sophisticated techniques based on integer programming, genetic algorithms or heuristics for the solution of the problem. In the Software as a Service area, energy consumption has not been considered yet. Further, a gap exists between the problem of maximization of QoS for end users and the optimum trade-off between performance and energy determined at the technological layer. The aim of this research is to fill this gap by privileging energy considerations and by implementing run-time optimum

allocation polices at the technological layer and IS-control modules. The adoption of integrated run-time mechanisms is advocated also in (Feng 2008). Self-* properties for systems as introduced by (Kephart 2003), provide the ability of creating Self-CHOP systems, i.e., Self Configuring, Self Healing, Self Optimizing, Self Protecting systems. Examples of applications are in the areas of Self Adaptive Middleware, Self Healing Databases, Autonomic Server Monitoring, and mobile Information Systems (Penici 2006). While currently this approach is focusing on middleware technology for infra-structure management, such an approach provides the basis for creating systems which change their way of working according to internal and external conditions, and set the basis for further investigation for a comprehensive approach towards the development of green IT for Information Systems.

Concerning data management, data deduplication (Elmagarmid 2007) is a method of reducing storage requirements by eliminating redundant data. In a backup and recovery system, gigabytes of information will be stored over and over again at each backup, creating several copies of data that might change only incrementally over time. With data deduplication, the basic idea is to store only data changes on storage devices, while redundant data is replaced with a pointer to the unique data copy. Current technology has applied data deduplication for the management of system backup enabling up to 300:1 data reduction for backup and storage applications (McGillicuddy 2007). The starting point for data dedupliction is the identification of the redundancies that is the recognition of representations of the same object in different sources (Newcombe 1962). This problem has been known for more than decades as the record linkage or the record matching problem. This research can contribute to a comprehensive green IT approach for Information Systems to support the identification of redundancies in stored files by considering the metadata that describe them and thus by considering semi-structured data. Anyway, data deduplication has to be performed by considering only relevant and useful data. The evaluation of data relevance is an open issue in the literature. Many times relevance is judged by considering subjective evaluations. Previous work just provides a set of guidelines based on a few of quantitative and qualitative parameters. The identification of significant qualitative and quantitative parameters for relevance assessment is needed, in particular considering time aspects for data objects that are not temporally valid or not frequently used.

3. Design of green Information Systems and service-based approaches

In order to evaluate the impact of use of energy and to design new energy efficient Information Systems, innovation is needed both at the strategic, governance, and technology levels (see figure below).

An integrated approach for energy management of systems supporting both applications and data is advocated. A service-oriented approach to application development offers a flexible support for application execution. In forthcoming service technology, adaptivity becomes a major asset, and sets the basis for developing flexible systems which are able to react to variable operation conditions (S-Cube Network, 2008).
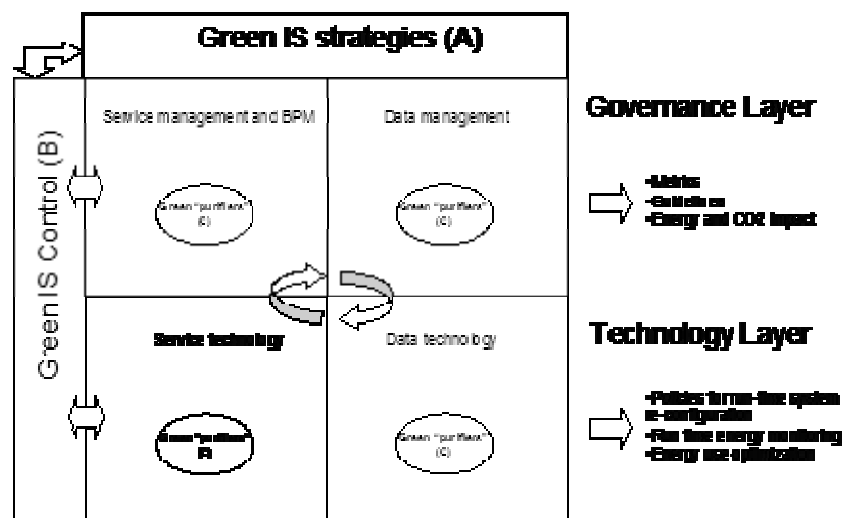

F**ig. 1.** An integrated approach to green Information Systems

At the strategic level, Green IS strategies have to be developed, considering the context of operation of the system, the main business goals, and the general characteristics of the underlying infrastructure. Such strategic level is needed to define the requirements and constraints for a green IS approach in IT systems.

The IT system provides the infrastructure for running the information system and can be considered as composed by a Governance layer and a Technology layer.

In considering services, we address both services provided inside organizations and external services, as well as the current trend towards out-sourcing and using software as a service.

While existing research is currently focusing on the infrastructure level for service technology, such as in computer centres, and on hardware technology, we propose the development of new original methods for reducing unwanted activities and data towards a "purifying" approach. The general idea is that of identifying improper use of resources in the system and sup-porting their reduction. The proposed approach can be based on the development of research in adaptive, self-healing, and self-managing systems to reduce the energy consumption, also considering external usage and environmental factors. In this sense, we propose the development of research activities towards the realization of

"green purifiers" for each of the four components (service governance, data governance, service technology, data technology).

In particular, at the service governance level, research has to focus on new techniques to analyze processes under a Green IT perspective. While Business Process Management technology focuses on process performance and mining, the focus of the research in this area will be (i) on the analysis of the process structure and intermediate data exchanges to reduce the generation of output which is not used within the process or in other processes, (ii) to analyze the energy consumption parameters to support process scheduling when resources are available or less energy is needed, instead of acquiring further computing resources for their execution.

A self-approach as proposed in the service engineering area (Kephart and Chess 2003) has to be investigated, in particular to adapt systems to variable operating conditions and to different phases of the life-cycle of running and executed applications and services. Frameworks for adapting service-based business processes such as the PAWS framework (Ardagna et al. 2007) can be adopted to realize context-aware and sustainable Information Systems.

However, such components cannot be used in isolation and have clear potential interrelations. The monitoring, analysis, and control of the use of such technological components have to be supported by research towards the development of a Green IS Control system, which will provide filtering and decision support functionalities to achieve the goals defined in the Green IS strategies. Research on Green IS strategies will focus on developing a comprehensive model to analyze the energy consumption impact of Information Systems, and for developing methodological guidelines to improve it.

4. Concluding Remarks

The present paper has presented a global approach to designing sustainable information systems, with a particular focus on service-based business processes.

A context-aware and "purifier-based" approach to Information Systems has been proposed, which has to be taken into consideration in the process design phases to result applicable during system operation.

References

1. S. Abiteboul. Issues in Monitoring Web Data. Proc. of the 13th Int. Conf. on Database and Expert Systems Applications, 2002, France. LNCS, 2453. Berlin: Springer, 1-8.

2. D. Ardagna, M. Comuzzi, E. Mussi, P. Plebani, B. Pernici. PAWS: a framework for processes with adaptive Web services. IEEE Software. 24(6), 39-46, 2007

3. D. Ardagna and B. Pernici. Adaptive service composition in flexible processes. IEEE Trans. on Software Engineering, June 2007.

4. D. Ardagna, M. Trubian, and L. Zhang. SLA based resource allocation policies in autonomic environments. Journal of Parallel and Distributed Computing, 67(3):259–270, 2007.

5. L. A. Barroso, U. Hölzle. The Case for Energy-Proportional Computing, IEEE Computer, vol. 40, 2007.

6. M. Borup. Green Technology Foresight as Instrument in Governance for Sustainability. Proceedings of the 2003 Berlin Conference on the Hu-man Dimensions of Global Environmental Change, 386 - 408.

7. G. Canfora, M. Penta, R. Esposito, and M. L. Villani. QoS-Aware Re-planning of Composite Web Services. In ICWS 2005 Proc., 2005. Orlando.

8. J. S. Chase and D. C. Anderson. Managing Energy and Server Re-sources in Hosting Centers. In ACM Symposium on Operating Systems principles, 2001.

9. N. Cohen. Greening the Internet: ten ways e-commerce could affect the environment and what we can do. Environ Qual. Manag. 9 (1), 1-16, 1998.

10. A. K. Elmagarmid, P. G. Ipeirotis, and V. S. Verykios. Duplicate record detection: A survey. IEEE Trans. Knowl. Data Eng., 19(1):1–16, 2007.

11. W. Feng, X. Feng, R. Ge. Green Supercomputing Comes of Age. IT Professional, 10(1), 17-23, 2008.

12. M. A. Hernandez, S. J. Stolfo. Real-world data is dirty: Data cleansing and the merge/purge a problem. Data Min. Knowl. Discov., 2(1):9–37, 1998.

13. Hodge, G.M. Best practises for digital archiving, D-Lib, January 2000.

14. P. Jokinen, P. Malaska, J. Kaivo-Oja. The environment in an "Informa-tion society": a transition stage towards more sustainable development? Futures 30 (6), 485-498, 1998.

15. J. O. Kephart, D. M. Chess. The Vision of Autonomic Computing. IEEE Computer, 36 (1), 41-50, 2003.

16. J. Kephart, H. Chan, R. Das, D. Levine, G. Tesauro, F. Rawson, and C. Lefurgy. Coordinating Multiple Autonomic Managers to Achieve Speci-fied Power-performance Tradeoffs. In ICAC Proc., June 2007.

17. U. Kremer, J. Hicks, J. M. Rehg. A Compilation Framework for Power and Energy Management on Mobile Computers. In LCPC Proc., 115-131, 2001.

18. D. Kusic and N. Kandasamy. Risk-Aware Limited Lookahead Control for Dynamic Resource Provisioning in Enterprise Computing Systems. In ICSOC 2004 Proc., 2004.

19. Z. Maamar, Q.Z. Sheng, B. Benatallah. Interleaving Web Services Composition and Execution Using Software Agents and Delegation. In WSABE Proc, 2003.

20. S. McGillicuddy. Data deduplication technology of growing interest to SMBs,http://searchcio-midmar-ket.techtarget.com/news/article/0,289142,sid183_gci1265307,00.html, 2007.

21. V. Metha. A Holistic Solution to the IT Energy Crisis. http://www.greenercomputing.com/reviews_third.cfm?NewsID=36321, 2007.

22. S. Murugesan, Harnessing green IT: Principles and practice. IEEE IT Professional 10(1), 24-33, 2008.

23. H.B. Newcombe and J.M. Kennedy. Record Linkage: Making Maxi-mum Use of the Discriminating Power of Identifying Information. Comm. of the ACM, 5 , 563-566, 1962.

24. B. Pernici (ed). Mobile Information Systems Infrastructure and Design for Adaptivity and Flexibility, Springer 2006

25. S-Cube Network, http://www.s-cube-network.eu/, 2008

26. J. Williams, Joseph, L. Curtis. Green: The New Computing Coat of Arms? IT Professional, 10(1), 12-16, Jan.-Feb. 2008.

27. L. Yi, H. R. Thomas. A review of research on the environmental impact of e-business and ICT. Environment International 33(6): 841-849, August 2007.

28. T. Yu, Y. Zhang, and K.-J. Lin. Efficient algorithms for web services selection with end-to-end QoS constraints. ACM Trans. Web, 1(1):1–26, 2007.

29. L. Zeng, B. Benatallah, M. Dumas, J. Kalagnamam, and H. Chang. QoS-aware middleware for Web services composition. IEEE Trans. on Software Engineering, 30(5), May 2004.