Employing Virtual Power Analytics and Linked Data for Enterprise IT Energy Informatics

Dippy Aggarwal
University of Cincinnati
2600 Clifton Ave
Cincinnati, OH, USA
aggarwdy@mail.uc.edu

Edward Curry
Insight Centre for Data Analytics
National University of Ireland
Galway, Ireland
edcurry@acm.org

Karen C. Davis
University of Cincinnati
2600 Clifton Ave
Cincinnati, OH, USA
karen.davis@uc.edu

ABSTRACT

The paper provides a solution to an organization's challenge for reducing its ecological footprint. We develop a framework for visualizing information concerning power consumption of IT devices in real-time using power metering approaches and leverage semantic web technologies to weave information that is dispersed across different data sources, ranging from an organization's asset databases to excel spreadsheets to the data available on manufacturer's websites. Visualizing the aforementioned information and bringing it together in one application can help organizations in understanding and adapting their energy consumption behavior in order to better support environmentally sustainable practices. The purpose of this project, from a power analytics standpoint, is twofold: (1) it aims to provide users with insight into their device's current power consumption without leaving a footprint on their machine, and (2) to analyze if one method gives more accurate results than others. We performed analysis using Microsoft Joulemeter and power estimation using software design methodologies (with numbers from smart meters as our benchmark).

Categories and Subject Descriptors

H.4 [**Information Systems Applications**]: Miscellaneous; D.2.8 [**Software Engineering**]: Metrics

General Terms

Experimentation, Management, Measurement, Performance

Keywords

Green IT, Semantic Web, Energy Informatics, Power Analytics, Data Integration, Linked Data, Semantic Web

1. INTRODUCTION

The field of energy informatics has gained significant attention in the past few years. This is evident from the various research efforts [5, 3, 4, 6] and industries[1] [4, 2] recognizing the business value in adopting environmentally sustainable business practices. Enterprise Computing[2] has emerged as an umbrella term encompassing the idea of observing business at a holistic level. This includes defining concepts and ideas that are common and pertinent across all the units/departments within an organization. In this context, enterprise energy informatics is concerned with designing and implementing information systems that can provide information on energy consumption within an enterprise.

In our work, we aim to capture the ecological footprint attributable to IT. Green IT is emerging as an equally important quality metric for an organization's market standing, as are the traditional goals of customer satisfaction and profitability. Studies on user's device usage patterns[13] have shown that a lack of awareness and incentive in terms of economic benefits serve to de-motivate users from demonstrating their environmental responsibility. As an example, many times people do not put their PC to sleep mode even when they know that they will not be working on it for few hours since they do not want to wait for it to power up again. As another example, many people have not activated power saving settings on their PC, although they come installed by default [13]. There is a lack of transparency in realizing the monetary gains an organization can achieve by saving on the energy consumption of their IT devices. Chheda et al. [1] have employed simple math to show that running just one server can cost a company $301.54 per year and produce environmental cost of $7,731.8$ pounds of CO_2.

Our project seeks a solution to an organization's challenges for reducing its ecological footprint, in two aspects, (a) power analytics challenge, and (b) information management challenge.The power analytics challenge involves providing visibility to the end-user about the power consumption of their device in real-time. The information management challenge strives to address data integration by providing a model for linking information concerning devices within and from outside the organization. The information is dispersed across several data sources ranging from Excel files, company's internal databases, and manufacturer's websites. Representing data in the linked data format allows its

[1]Apple and the environment. Available at: `http://www.apple.com/environment/`
[2]Enterprise computing. Available at: `http://www.peterindia.net/EnterpriseComputOverview.html`

easy integration with data from other relevant sources[3] [8].

The rest of the paper is organized as follows. In Section 2 we enumerate our research contributions. Section 3 discusses the background and related work providing an overview of semantic web technologies and three power metering approaches. In Section 4 we present our approach and Section 5 concludes the paper with directions for future work.

2. RESEARCH CONTRIBUTIONS

The project makes the following two contributions: (a) perform a comparative study of two existing software-based power metering solutions, namely, Microsoft Joulemeter [10] and the power estimation modeling approach laid out by Lien et al. [11], and (b) develop a user interface to provide a holistic view of IT energy management within an enterprise. This involves bringing together all the related or overlapping information about a device dispersed across different data sources within an organization and web together in one place. We leverage linked data principles to solve the data integration challenge.

3. BACKGROUND AND RELATED WORK

This section, provides an overview of semantic web technologies and the three power metering approaches that we have studied in our project.

3.1 Power Analytics

Currently there are myriad of solutions available in the market that provide numbers for power consumption of a device in real-time. We perform a comparative study amongst them to determine advantages and limitations of one over the other.

A smart meter is wireless device that measures the power consumption of a device attached to it. The monitored device's power cord is connected to the meter that is plugged into external power socket. The power consumption readings are transmitted wirelessly over a USB modem connected to the device itself. However, it is not feasible from an economical as well as a mobility standpoint to have a smart meter for each device in an organization. Thus, it becomes imperative to leverage software-based models for power estimation.

The Microsoft Joulemeter is a software tool [10] that estimates the power consumption of a PC, virtual machine or software system. In our project we focused only on physical devices, laptops in particular. This solution suffers from the limitation of the need for installation on every machine that needs to be monitored and it specifically requires the Windows 7 operating system.

The approach, proposed by Lien et al. [11] estimates the server's power consumption at a given point of time, based on four parameters, (a) average power consumption of the device in idle mode (when CPU utilization is 0%), (b) average power consumption when CPU utilization is 100%, (c) CPU utilization at the time of measurement, and (d) constant β that is determined from experiments. Lien at al. [11] employ an equation for power modeling of streaming media servers. Our project leverages their equation model for laptops. The equation is modeled as:

$$P = D + (M - D)U^{1/\beta}$$

where

P represents the estimated power consumption,
D represents average power consumption in idle mode,
β represents, constant determined from experiments,
U represents CPU utilization, and
$M - D$ represents average power consumption in addition to D.

3.2 Information Management Challenge

Management of any IT infrastructure benefits from multiple data sources, ranging from an organization's internal databases, Excel spreadsheets and data available on the web [7]. With the current relational database technology, it can be argued that all this information can be brought together and stored in a relational database. However, semantic web technologies tend to be a stronger candidate solution for this scenario. The semantic web offers flexibility by removing the requirement for a pre-defined schema and enabling the addition of information in data stores as and when it becomes available.

The semantic web[4] is a widely-embraced paradigm in which semantic information in the form of a URI (Uniform Resource Identifier) is attached to data to provide a unique identity for it and to allow it to be processable by a human user as well as machine. This should not be misconstrued as a technique by which a computer will be able to comprehend the meaning of the data, but with the data tagged with semantic labels, a machine will be able to connect and process them, with relationships defined over them such as *sameAs* and *composedOf*.

Linked Data is a technology that is built upon current Web architecture and open standards proposed for semantic web [6]. The core idea behind this technology is to provide support towards large scale integration and reasoning over data. Resource Description Framework forms the standard format to represent linked data. It follows a graph based paradigm with the labels of the graph describing the data. RDF facts are specified as statements which are expressed in the form of a triple: subject, predicate and object [6].

4. APPROACH

We follow a two-step approach for power metering. In the first step, we perform a comparative study of two software-based approaches. In the second step, we visualize power consumption of a device on a dashboard in real-time. The architecture as shown in Figure 1 consists of four independent components: (1) device (laptop) that is being monitored for power consumption, (2) the software application-Microsoft Joulemeter that is installed on the laptop, (3) smart meter plugged into external power socket and. (4) a command based utility - typeperf[5] that will ping the laptop to record energy consumption readings periodically.

[3] Green and sustainable IT at DERI. available at: `http://dgsit.deri.ie/`

[4] An introduction to tim berners-lee's semantic web. available at:`http://www.techrepublic.com/article/an-introduction-to-tim-berners-lees-semantic-web`

[5] Typeperf. available at:`http://technet.microsoft.com/en-us/library/bb490960.aspx`

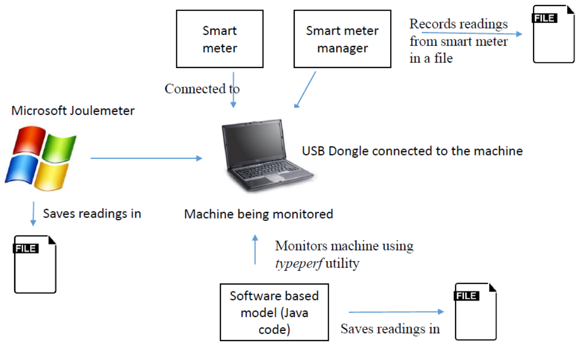

Figure 1: Architecture for power metering approach

The values recorded by each of three power metering approaches are saved in separate files. Figure 2 depicts the power consumption readings for a DELL laptop using three different approaches, namely smart meter, Joulemeter and software model by Lien et al. [11].

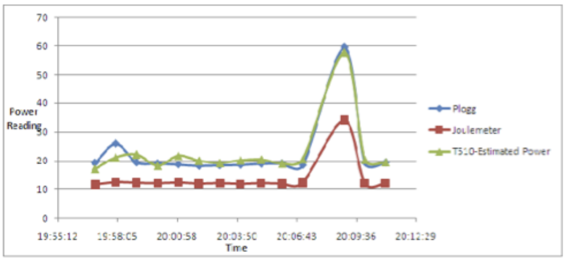

Figure 2: Power Consumption readings for DELL-T510 (without batteries)

In another experiment, as shown in Figure 3 we compared readings from software based model [11] with the readings from the smart meter. This experiment was performed over a different model of laptop, namely DELL-T61.

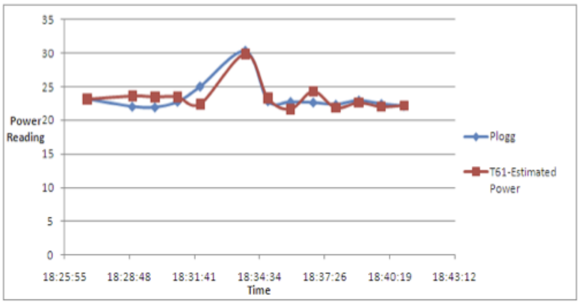

Figure 3: Power Consumption readings for DELL-T61 (without batteries)

The initial experiments show encouraging results for the accuracy of the readings obtained from using the model proposed by Lien at el. [11] as opposed to readings from Joulemeter. The values obtained from the smart meter were used as benchmark to compare the two sets of readings.

For the information management challenge, we leverage semantic data principles by (a) querying the organization's data, stored on a RDF server, and (b) as a tool to build linked data, by pulling out extra information about a laptop, which is not available in organization's data stores, from the data available on the web. The papers [6, 9, 14] describe the approach and ontology in detail for adapting the raw readings coming from the tools such as typerperf or sensors to linked data format. Figure 4 shows an overview of the ontology used in our work.

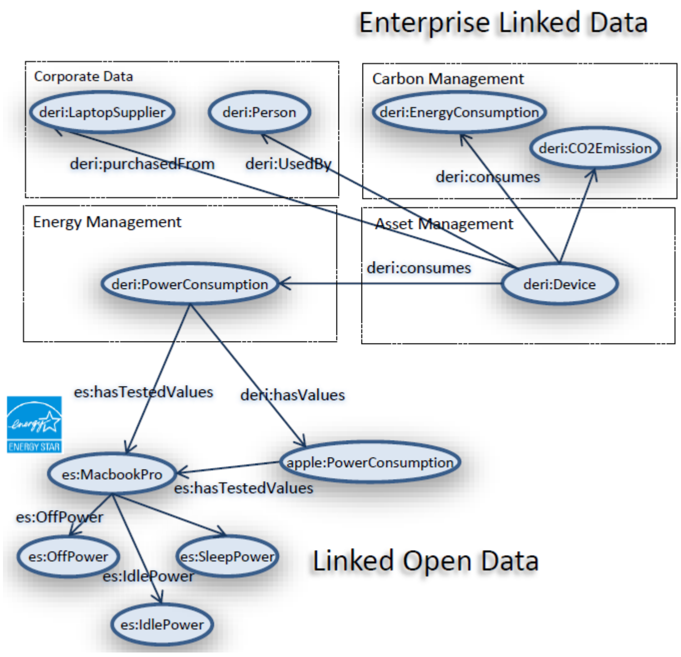

Figure 4: Ontology for IT system management

The ontology in Figure 4 describes IT devices, represented by node *deri:Device*, and the information about their power consumption, manufacturer and the person who owns the device. The two properties, *es:hasTestedValues* and *deri: hasValues* represent the information on energy consumption of a machine provided by two different sources, namely the manufacturer's website and Environmental Protection Agencies like Energy Star.

Figure 5 shows the dashboard providing a unified view of IT devices (laptops) built using the Linked Energy Intelligence Platform [6].

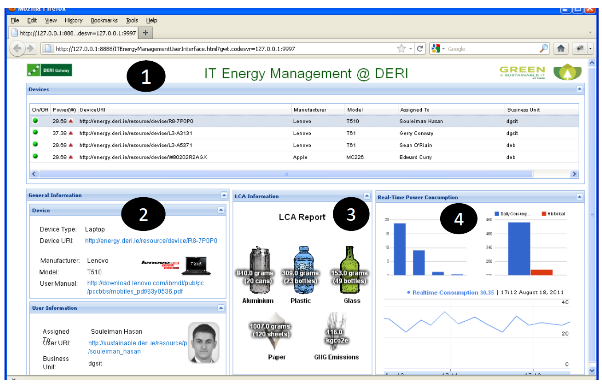

Figure 5: Unified view of IT energy management

Figure 5 can be analyzed in four sections, circled one to four. Sections 1 and 2 depict device's information pulled from organization's internal RDF server. Section 3 visualizes the lifecycle assessment information and Section 4 depicts the power consumption of laptop selected in Section 1. The life cycle assessment model measures the environmental effect of a device based on data from the manufacturer's supply chain and the device's power consumption from the use phase and through recycling. The environmental effect of devices can be quantified in terms of their material usage and energy efficiency[1]. This data could be published by the manufacturer as Linked Open Data on their website. The top-left bar graph in Section 4 shows the power consumption in different modes: active, idle, hibernate. The bar chart on top-right in Section 4 aims to provide visualization of current power consumption of device with respect to its historical (previous month) consumption and the line graph shows the power consumption of device in real-time. There is substantial potential for Energy Information Systems [14] for empowering organizations to improve their contribution towards environmental sustainability by developing energy management applications including energy tracking dashboards, to facilitate the understanding of energy consumption and environmental impact analysis of the organization as a whole.

5. CONCLUSION

The main contribution of the project is to provide visibility to the user into the information concerning their device usage that can help them to identify policies for achieving sustainability objectives. We perform a comparative study of two software-based solutions for power metering, Microsoft Joulemeter and Power Estimation Modeling [11], using the power readings obtained from smart meters as a benchmark. Empirical results from our initial set of experiments show strong results for the model proposed by Lien et al. [11], in terms of its ability to remotely estimate power consumption of a Windows machine, removing the requirement of installing software on every monitored machine without compromising on the accuracy level in the results. We also developed a user interface to depict power consumption of devices in real-time and to demonstrate the potential of leveraging semantic web technologies for addressing data integration challenges.

The future work will include conducting additional experiments to validate the strength of our results. Some ideas about potential additional experiments include obtaining empirical results for different machines for the same manufacturer-model combination. Thus, instead of generalizing results for a Dell-T510 based on experiments conducted on just one machine, it would provide a much higher degree of confidence if we have similar results for the same manufacturer-model combination. Secondly, we envision great potential for software based energy profiling for IT devices and systems other than laptops such as printers and data centers. There has been related work in the area [12], but we believe that it would be fruitful to survey the existing solutions and draw a comparison analysis chart. Another direction of research would be to capture and visualize energy consumption of multiple laptops at the same time, so that an energy profile can be observed at a team level or for a specific group of users within an organization.

6. REFERENCES

[1] Chheda, R., Shookowsky, D., Stefanovich, S., and Toscano, J. Profiling energy usage for efficient consumption. *The Architecture Journal: Green Computing Issue* (2008).

[2] Curry, E., and Donnellan, B. Implementing sustainable it strategy: the case of intel. *Journal of Information Technology Teaching Cases 4*, 1 (2014), 41–48.

[3] Curry, E., Guyon, B., Sheridan, C., and Donnellan, B. Developing a sustainable it capability: Lessons from intel's journey. *MIS Quarterly Executive 11*, 2 (2012).

[4] Curry, E., Guyon, B., Sheridan, C., and Donnellan, B. Sustainable it: Challenges, postures, and outcomes. *Computer 45*, 11 (2012), 0079–81.

[5] Curry, E., Hasan, S., Herstand, M., O'Riain, S., et al. An entity-centric approach to green information systems. In *Proceedings of the 19th European Conference on Information Systems (ECIS 2011)*.

[6] Curry, E., Hasan, S., and O'Riain, S. Enterprise energy management using a linked dataspace for energy intelligence. In *Sustainable Internet and ICT for Sustainability (SustainIT), 2012* (2012), IEEE.

[7] Feridun, M., and Tanner, A. Using linked data for systems management. In *Network Operations and Management Symposium (NOMS), 2010 IEEE* (2010), IEEE, pp. 926–929.

[8] Freitas, A., Curry, E., Oliveira, J. G., and O'Riain, S. Querying heterogeneous datasets on the linked data web: Challenges, approaches, and trends. *Internet Computing, IEEE 16*, 1 (2012), 24–33.

[9] Hasan, S., Curry, E., Banduk, M., and O'Riain, S. Toward situation awareness for the semantic sensor web: Complex event processing with dynamic linked data enrichment. *Semantic Sensor Networks* (2011).

[10] Kansal, A., Zhao, F., Liu, J., Kothari, N., and Bhattacharya, A. A. Virtual machine power metering and provisioning. In *Proceedings of the 1st ACM symposium on Cloud computing* (2010), ACM, pp. 39–50.

[11] Lien, C.-H., Liu, M. F., Bai, Y.-W., Lin, C. H., and Lin, M.-B. Measurement by the software design for the power consumption of streaming media servers. In *Instrumentation and Measurement Technology Conference, 2006. IMTC 2006. Proceedings of the IEEE* (2006), IEEE, pp. 1597–1602.

[12] Lim, Harold and Kansal, Aman and Liu, Jie Power budgeting for virtualized data centers. In *USENIX Annual Technical Conference (USENIX ATC'11)/* (2011), pp. 59.

[13] Mattern, F., Staake, T., and Weiss, M. Ict for green: how computers can help us to conserve energy. In *Proceedings of the 1st international conference on energy-efficient computing and networking* (2010), ACM, pp. 1–10.

[14] Curry, E., and Donnellan, B. Sustainable information systems and green metrics. *Harnessing Green IT: Principles and Practices, John Wiley & Sons, Inc*, (2012), 167–198.