Leveraging Renewable Energy in Data Centers: Present and Future

Keynote Summary

Ricardo Bianchini
Department of Computer Science
Rutgers University
ricardob@cs.rutgers.edu

ABSTRACT

Interest has been growing in powering data centers (at least partially) with renewable or "green" sources of energy, such as solar or wind. However, it is challenging to use these sources because, unlike the "brown" (carbon-intensive) energy drawn from the electrical grid, they are not always available. In this keynote talk, I will first discuss the tradeoffs involved in leveraging green energy today and the prospects for the future. I will then discuss the main research challenges and questions involved in managing the use of green energy in data centers. Next, I will describe some of the software and hardware that researchers are building to explore these challenges and questions. Specifically, I will overview systems that match a data center's computational workload to the green energy supply. I will also describe Parasol, the solar-powered micro-data center we have just built at Rutgers University. Finally, I will discuss some potential avenues for future research on this topic.

Categories and Subject Descriptors

A.1 [**Introductory and Survey**]; C.5.5 [**Computer System Implementation**]: Servers; D.4.1 [**Operating Systems**]: Process Management

Keywords

Renewable energy, energy-aware scheduling, data centers.

1. INTRODUCTION

Data centers consume an enormous amount of energy [13]. In recent years, large data center operators, like Google and Microsoft, have significantly improved the energy efficiency of their multi-megawatt data centers. However, the majority of the energy consumed by data centers is actually due to small and medium-sized data centers [13], which are much more numerous and much less efficient. These facilities may range from a few dozen servers housed in a machine room to several hundreds of servers housed in a larger enterprise installation. Many of these facilities run high-performance computing workloads, such as data analytics and scientific simulations.

The energy consumed by data centers represents a financial burden on the organizations that operate them, and an infrastructure burden on power utilities. In addition, data centers contribute to climate change, since most of the electricity produced in the US and around the world derives from carbon-intensive fuels, such as coal and natural gas.

HPDC'12, June 18–22, 2012, Delft, The Netherlands.
ACM 978-1-4503-0805-2/12/06.

Due to these concerns and societal pressure, we are starting to see many "green" data centers powered (at least partially) by renewable sources of energy [2, 4]. These green data centers either generate their own electricity or draw directly from a nearby renewable power plant. Among other advantages, these types of renewable plant/data center co-location reduce the energy losses involved in power conversion and transmission over long distances.

Solar and wind are two of the most promising sources of green energy for data centers, as they are clean and broadly available. However, solar and wind do have two main limitations today: the space they require and their capital costs.

Fortunately, predicted improvements in efficiency and reductions in cost/Watt will alleviate these problems significantly in the future. For example, improvements in photovoltaic (PV) solar panels and new PV technologies are expected to triple today's efficiencies until 2030 [12]. Over the same period, the cost/Watt of PV panels is expected to become less than half of what it is today. In addition, governments currently provide incentives for green energy generation. For example, federal and state incentives in New Jersey can reduce the capital cost of solar installations by 60% [3]. If these incentives continue, cost may not be a significant factor in the future.

These trends suggest that solar and/or wind power will become increasingly attractive, especially for small and medium data centers as they require smaller and cheaper installations. Moreover, solar panels and/or wind turbines can be deployed in small increments for these data centers.

2. RESEARCH CHALLENGES

The main challenge with solar and wind energy is that, unlike brown energy drawn from the grid, it is variable. To mitigate this variability, data centers could "bank" green energy in batteries or on the grid itself. However, these approaches incur energy losses and high additional costs in the case of batteries. Instead, data centers can maximize the use of the available green energy by matching the energy demand (computational work) to the supply.

The need to match energy demand and supply prompts many interesting research questions. For example, what kinds of data center workloads are amenable to green data centers? What kinds of techniques can we apply to better match the demand for energy to the variable energy supply? Should we allow programmers to specify what types of techniques can be used? How well can we predict solar and wind availability? If batteries are available, how should we manage them? Can we leverage geographical distribution to maximize our use of green energy? If we have a choice, where should be place green data centers to strike a good compromise between high energy generation and data center costs?

Researchers have started building software and hardware to address these questions, e.g. [5, 6, 7, 8, 9, 10, 11]. The next section describes some software efforts, whereas Section 4 describes Parasol, a solar-powered μdata center we just built at Rutgers.

3. SOFTWARE FOR GREEN DATA CENTERS

We have recently built two load-scheduling systems for green data centers: GreenSlot [5] and GreenHadoop [6]. Both systems assume that (1) the data center is connected to a solar array and the electrical grid, and (2) there are no batteries. Their goal is to maximize the use of solar energy; brown energy should only be consumed when solar energy is not available.

GreenSlot extends the SLURM scheduler for batch jobs. GreenSlot maximizes the solar energy consumption while meeting the jobs' deadlines. If brown energy must be used to avoid deadline violations, GreenSlot schedules jobs for times when brown energy is cheap. In more detail, it first predicts the amount of solar energy that will likely be available in the future, using historical data and weather forecasts. Based on its predictions and the information provided by users, it schedules the workload by creating resource reservations into the future. Whenever servers are not needed, GreenSlot transitions them to a sleep (ACPI S3) state.

Along similar lines, GreenHadoop extends the Hadoop data-processing framework. Scheduling the energy consumption of Hadoop jobs is challenging, because they do not specify the number of servers to use, their run times, or their energy needs. Moreover, power-managing servers here requires guaranteeing that the data to be accessed by the active jobs remains available. Besides managing energy consumption and brown energy costs, GreenHadoop manages the cost of peak brown power consumption.

Instead of adapting to the availability of green energy through job scheduling, researchers from UMass Amherst have proposed to modulate the servers' duty cycle using fast sleep states [10]. In contrast, researchers from UC Berkeley have proposed to adjust the quality of the replies provided to users in interactive workloads [7]. For mixed interactive and batch workloads, researchers from UC San Diego have proposed to adapt the amount of batch processing dynamically [1]. Finally, several groups have considered load distribution across green data centers to "follow the renewables", e.g. [8, 9, 11].

4. PARASOL: OUR GREEN DATA CENTER

Parasol is our research platform for studying the use of renewable energy in data centers. It comprises a small container, a set of solar panels, and batteries. The container lies on a steel structure placed on the roof of our building. The 16 solar panels are mounted on top of the steel structure and shade the container from the sun most of the time. We expect that the panels will produce up to 3KW of AC power (after derating). Figure 1 shows the steel structure, the container, and the solar panels.

The container hosts two racks of energy-efficient IT equipment. The racks currently host 64 Atom-based half-U servers equipped with solid-state drives, but we will install many more of these servers soon (the maximum capacity of Parasol is roughly 150 of these servers). The container uses free cooling whenever possible, and direct-exchange air conditioning (HVAC) otherwise. Our desire to study free cooling is the main reason we place the servers on the roof, rather than inside our building. Besides the solar panels, Parasol can draw energy from its batteries and/or the electrical grid. Three manual switches enable different configurations for the supply of energy. For example, we can configure Parasol to operate completely off the electrical grid.

Parasol includes an extensive monitoring infrastructure to quantify resource utilization, power generation and consumption, server and data center temperatures, and air flow and quality. A powerful server located in one of the racks collects all the monitoring information and backs it up to our main laboratory.

Figure 1: Final stage of the construction of Parasol.

Acknowledgements

I am very grateful to my collaborators in this area: Josep L. Berral, Íñigo Goiri, Jordi Guitart, Md E. Haque, William Katsak, Kien Le, Margaret Martonosi, Thu D. Nguyen, and Jordi Torres. My group's research has been partially funded by NSF grants CSR-0916518 and CSR-1117368, and the Rutgers Green Computing Initiative.

5. REFERENCES

[1] B. Aksanli et al. Utilizing Green Energy Prediction to Schedule Mixed Batch and Service Jobs in Data Centers. In *HotPower*, 2011.

[2] Data Center Knowledge. Apple Plans 20MW of Solar Power for iDataCenter, 2012. http://www.datacenterknowledge.com/archives/2012/02/20/-apple-plans-20mw-of-solar-power-for-idatacenter/.

[3] DSIRE. Database of State Incentives for Renewables and Efficiency. http://www.dsireusa.org/.

[4] EcobusinessLinks. Green Hosting - Sustainable Solar & Wind Energy Web Hosting, 2012. http://www.ecobusinesslinks.com/green_web_hosting.htm.

[5] I. Goiri et al. GreenSlot: Scheduling Energy Consumption in Green Datacenters. In *Supercomputing*, 2011.

[6] I. Goiri et al. GreenHadoop: Leveraging Green Energy in Data-Processing Frameworks. In *EuroSys*, 2012.

[7] A. Krioukov et al. Design and Evaluation of an Energy Agile Computing Cluster. Technical Report EECS-2012-13, UC Berkeley, 2012.

[8] K. Le et al. Cost- And Energy-Aware Load Distribution Across Data Centers. In *HotPower*, 2009.

[9] Z. Liu et al. Greening Geographical Load Balancing. In *SIGMETRICS*, 2011.

[10] N. Sharma et al. Blink: Managing Server Clusters on Intermittent Power. In *ASPLOS*, 2011.

[11] C. Stewart and K. Shen. Some Joules Are More Precious Than Others: Managing Renewable Energy in the Datacenter. In *HotPower*, 2009.

[12] Technology Roadmap – Solar Photovoltaic Energy. International Energy Agency, 2010. http://www.iea.org.

[13] US Environmental Protection Agency. Report to Congress on Server and Data Center Energy Efficiency, 2007.