Ad Hoc Networks 25 (2015) 520–534

Contents lists available at ScienceDirect

Ad Hoc Networks

journal homepage: www.elsevier.com/locate/adhoc

Matching renewable energy supply and demand in green datacenters ☆

Íñigo Goiri [a], Md E. Haque [a,*], Kien Le [a], Ryan Beauchea [a], Thu D. Nguyen [a], Jordi Guitart [b], Jordi Torres [b], Ricardo Bianchini [a]

[a] *Department of Computer Science, Rutgers University, USA*
[b] *Universitat Politecnica de Catalunya and Barcelona Supercomputing Center, Spain*

ARTICLE INFO

Article history:
Received 31 March 2014
Received in revised form 22 September 2014
Accepted 10 November 2014
Available online 18 November 2014

Keywords:
Green energy
Energy-aware job scheduling
Datacenters

ABSTRACT

In this paper, we propose GreenSlot, a scheduler for parallel batch jobs in a datacenter powered by a photovoltaic solar array and the electrical grid (as a backup). GreenSlot predicts the amount of solar energy that will be available in the near future, and schedules the workload to maximize the green energy consumption while meeting the jobs' deadlines. If grid energy must be used to avoid deadline violations, the scheduler selects times when it is cheap. Evaluation results show that GreenSlot can increase solar energy consumption by up to 117% and decrease energy cost by up to 39%, compared to conventional schedulers, when scheduling three scientific workloads and a data processing workload. Based on these positive results, we conclude that green datacenters and green-energy-aware scheduling can have a significant role in building a more sustainable IT ecosystem.

1. Introduction

Datacenters consume an enormous amount of energy: estimates for 2010 indicate that they consume around 1.5% of the total electricity used world-wide [1]. Electricity cost thus represents a significant burden for datacenter operators. Moreover, this electricity consumption contributes to climate change, since most of the electricity is produced by burning fossil fuels. A 2008 study estimated world-wide datacenters to emit 116 million metric tons of carbon, slightly more than the entire country of Nigeria [2]. We refer to the energy produced by carbon-intensive means and distributed via the electrical grid as "brown energy".

These cost and environmental concerns have been prompting many "green" energy initiatives. One initiative is for datacenters to either generate their own renewable energy or draw power directly from a nearby renewable power plant. This approach is being implemented by many small and medium datacenters (partially or completely) powered by solar and/or wind energy all over the globe [3]. Larger companies are also investing in this direction. For example, Apple is building a 40 MW solar array for its North Carolina datacenter [4]. McGraw-Hill has recently completed a 14 MW solar array for its datacenter [5].

We expect that this trend will continue, as these technologies' capital costs keep decreasing (e.g., the inflation-adjusted cost of solar panels has decreased by 10-fold in the last three decades [6]) and governments continue to provide generous incentives for green power generation (e.g., federal and state incentives for solar power in the United States can reduce capital costs by up to 60% [7]). In fact, the trend may actually accelerate if carbon taxes and/or cap-and-trade schemes spread from Europe and

☆ This submission is a modified and extended version of "GreenSlot: Scheduling Energy Consumption in Green Datacenters", which was originally published in SC'11.

* Corresponding author.
E-mail address: mdhaque@cs.rutgers.edu (M.E. Haque).

http://dx.doi.org/10.1016/j.adhoc.2014.11.012
1570-8705/© 2014 Published by Elsevier B.V.

Asia to the rest of the world. For example, a cap-and-trade scheme in the UK imposes caps on the brown energy consumption of large consumers [8]. We present a more extensive discussion of the feasibility of using green energy in datacenters in [9].

We argue that the ideal design for green datacenters connects them to both the solar/wind energy source and the electrical grid (as a backup). The major research challenge with solar and wind energy is that, differently from brown energy drawn from the grid, it is not always available. For example, photovoltaic (PV) solar energy is only available during the day and the amount produced depends on the weather and the season. Datacenters sometimes can "bank" green energy in batteries or on the grid itself (called net metering) to mitigate this variability. However, both batteries and net metering have problems: (1) batteries involve energy losses due to internal resistance and self-discharge; (2) the cost of purchasing and maintaining batteries can dominate in a solar system [9,10]; (3) today's most popular battery technology for datacenters (lead-acid) uses chemicals that are harmful to the environment; (4) net metering incurs energy losses due to the voltage transformation involved in feeding the green energy into the grid; (5) net metering is not available in many parts of the world; and (6) where net metering is available, the power company may pay less than the retail electricity price for the green energy.

Thus, in this paper, we investigate how to manage a datacenter's computational workload to match the green energy supply. In particular, we design a scheduler for parallel batch jobs, called GreenSlot, in a datacenter powered by an array of PV solar panels and the electrical grid. Jobs submitted to GreenSlot come with user-specified numbers of nodes, expected running times, and deadlines by which they shall have completed. The deadline information provides the flexibility that GreenSlot needs to manage energy consumption aggressively.

GreenSlot seeks to maximize the green energy consumption (or equivalently to minimize the brown energy consumption) while meeting the jobs' deadlines. If brown energy must be used to avoid deadline violations, it schedules jobs for times when brown energy is cheap. In more detail, GreenSlot combines solar energy prediction, energy-cost-awareness, and least slack time first (LSTF) job ordering [11]. 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. When a job's scheduled start time arrives, GreenSlot dispatches it for execution. Clearly, GreenSlot differs significantly from most job schedulers, which seek to reduce completion times or bounded slowdown.

We implement two versions of GreenSlot: one extends the SLURM scheduler for Linux [12], and the second extends the MapReduce scheduler of Hadoop [13]. We use real scientific workloads from the Life Sciences Department of the Barcelona Supercomputing Center to evaluate our SLURM extension and a Facebook-inspired workload to evaluate our Hadoop extension. Our results show that GreenSlot accurately predicts the amount of solar energy to become available. The results also show that GreenSlot can increase green energy consumption and decrease energy cost by up to 117% and 39%, respectively, for the workloads/systems evaluated.

Based on these positive results, we conclude that green datacenters and green-energy-aware scheduling can have a significant role in building a more sustainable Information Technology ecosystem.

In summary, we make the following contributions: (1) introduce GreenSlot, a batch job scheduler for datacenters partly powered by solar energy; (2) implement and evaluate GreenSlot in two different environments: a scientific computing cluster and a data-processing MapReduce cluster; and (3) present extensive results isolating the impact of different aspects of the scheduler.

2. Background

Solar energy and datacenters. Solar is a promising clean energy technology, as it does not cause the environmental disruption of hydroelectric energy and does not have the waste storage problem of nuclear energy. Wind energy is also promising, but is not as abundant in many locations. Except for our (solar) energy predictions, our work is directly applicable to wind energy as well.

Transforming solar energy into (direct-current or DC) electricity is commonly done using PV panels. The panels are made of cells containing PV materials, such as monocrystalline and polycrystalline silicon. The photons of sunlight transfer energy to the electrons in the material. This energy causes the electrons to transfer between the two regions of the material, producing a current that is driven through the electrical load (e.g., a datacenter).

There are multiple ways to connect solar panels to a datacenter. Fig. 1 shows an example. The AC Load is the server and cooling equipment, which typically runs on alternating-current (AC) electricity. The DC electricity is converted to AC using an inverter. Excess solar energy can be stored in batteries via a charge controller. The controller may also connect to the electrical grid, in case the datacenter must operate even when solar energy is not available. Where net metering is available, one can feed excess solar energy into the grid for a reduction in brown energy costs.

The design we study in this paper does not include batteries or net metering, for the reasons we mentioned in the Introduction. We assume that the datacenter can be fully powered by the grid when insufficient green energy is being produced. On the other hand, any green energy that is not immediately used by the datacenter is wasted. Fortunately, GreenSlot is very successful at limiting waste. In fact, assuming the results from Section 5 and the best governmental incentives in the United States, the current

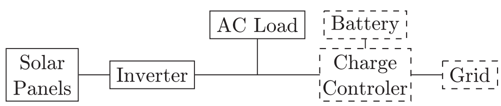

Fig. 1. Components of a solar-powered system. Dashed boxes represent optional components.