Future Generation Computer Systems 28 (2012) 405–414

Contents lists available at SciVerse ScienceDirect

Future Generation Computer Systems

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

Distributed computing for carbon footprint reduction by exploiting low-footprint energy availability

Ward Van Heddeghem *, Willem Vereecken, Didier Colle, Mario Pickavet, Piet Demeester

Department of Information Technology (INTEC) of Ghent University - IBBT, Gaston Crommelaan 8, B-9050 Gent, Belgium

ARTICLE INFO

Article history:
Received 29 October 2010
Received in revised form
11 April 2011
Accepted 7 May 2011
Available online 15 May 2011

Keywords:
Green ICT
Energy-efficiency
Power consumption
Distributed computing
Grid computing
Carbon footprint

ABSTRACT

Low carbon footprint energy sources such as solar and wind power typically suffer from unpredictable or limited availability. By globally distributing a number of these renewable sources, these effects can largely be compensated for. We look at the feasibility of this approach for powering already distributed data centers in order to operate at a reduced total carbon footprint. From our study we show that carbon footprint reductions are possible, but that these are highly dependent on the approach and parameters involved. Especially the manufacturing footprint and the geographical region are critical parameters to consider. Deploying additional data centers can help in reducing the total carbon footprint, but substantial reductions can be achieved when data centers with nominal capacity well below maximum capacity redistribute processing to sites based on renewable energy availability.

1. Introduction

Data center power consumption is significant, and growing—The last decade has seen a steady rise in data center capacity and associated power consumption. In 2008, the yearly average worldwide data center power consumption was estimated to be around 29 GW [1]. This is comparable to the total electricity consumption of Spain in the same year [2], a country that ranks in the top 15 of the list of electricity consumption per country. In [3], it was estimated that the aggregate electricity use for servers worldwide doubled over the period 2000–2005. With the predicted growth of Internet-based services for social networks and video, and with the growing usage of mobile thin clients such as smart phones that require a server back-end [4], it seems unlikely that this increase will halt soon.

Using renewable energy, in addition to energy-efficiency, is key to mitigate climate change—While the growing energy consumption in data centers presents some issues both economically and technically, there has been a growing concern from an environmental point of view as well, with electricity consumption contributing to greenhouse gas (GHG) emission. Two high-level approaches can help in reducing GHG emissions: (a) an improvement in energy-efficiency to reduce the amount of electrical energy used, and (b) use of energy that contributes little to GHG emissions. What concerns the latter, this electrical energy will typically come from renewable energy sources such as solar and wind power.

Adding renewable energy to the current energy mix still poses some issues—While renewable energy is indeed already promoted and used to mitigate climate change both in ICT and non-ICT sectors, significantly increasing the amount of renewable energy as part of the regular energy mix raises a number of issues [5]. First, because most good sites for renewable energy sources may be located in distant areas with limited transmission capacity, and it might take many years for the required transmission infrastructure to become available [6]. Second, the distributed power generation poses many challenges for the existing distribution infrastructure, especially with respect to protection and control strategies due to new flow patterns [6,7]. Third, with renewable energy sources likely to be located in distant areas, the transmission losses will increase; current transmission losses are already estimated to be around 6.5% of the total electricity disposition[1] for the USA in 2007 [8]. Fourth, with hydro-power usually reserved for peak power handling [9], other renewable energy sources such as wind and solar power are usually characterized by intermittent power delivery, resulting in periods of peak power being available and no power being available at all.

[1] To be correct, the loss percentage is calculated as a fraction of the total electricity disposition excluding direct use. Direct use electricity is electricity that is generated at facilities that is not put onto the electricity transmission and distribution grid, and therefore does not contribute to transmission and distribution losses [8].

* Corresponding author. Tel.: +32 0 9 33 14 977; fax: +32 0 9 33 14 899.
E-mail address: ward.vanheddeghem@intec.ugent.be (W. Van Heddeghem).

0167-739X/$ – see front matter © 2011 Elsevier B.V. All rights reserved.
doi:10.1016/j.future.2011.05.004

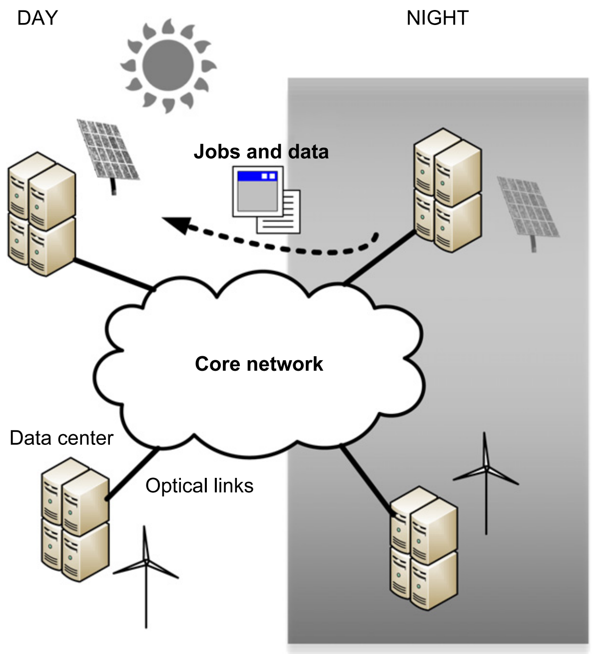

Fig. 1. Distributed data center.

Data centers are uniquely positioned to provide an alternative solution—Data centers have become more and more globally distributed for a number of reasons as summarized by [10]: "the need for high availability and disaster tolerance, the sheer size of their computational infrastructure and/or the desire to provide uniform access times to the infrastructure from widely distributed client sites". This geographical distribution of data centers, combined with the availability of low-power and high-speed optical links, allows them to be located near renewable energy sites. With technology currently available to migrate live virtual machines while minimizing or avoiding downtime altogether [11–13], jobs can be dynamically moved from a data center site where renewable power dwindles to a different site with readily available renewable power. This approach has previously been referred to as 'Follow The Sun/Follow The Wind' (FTSFTW) [5].

Fig. 1 illustrates this concept with solar powered data center sites. As the sun sets in the top-right data center (and the capacity of potential backup-batteries fall below a critical value) the site's data and jobs are moved to a different site (top left) where solar power has become available.

In this paper we will evaluate the carbon footprint and potential footprint savings of such a FTSFTW-based distributed data center. We will generalize on the notion of renewable energy, and instead consider low-footprint (LF) energy and high-footprint (HF) energy. As a metric for the carbon footprint we will use grams of CO_2-eq, unless otherwise indicated. CO_2-eq indicates CO_2-equivalent emissions, which is the amount of CO_2 that would have the same global warming potential when measured over a given time horizon (generally 100 years), as an emitted amount of a long-lived GHG or a mixture of GHGs.

The contributions of this paper are the following:

- we provide a mathematical model for calculating the carbon footprint and savings of such a distributed data center infrastructure which is powered by a fixed mix of LF and HF energy (Section 3),
- we provide a detailed and realistic quantification of the parameters in our mathematical formulation (Section 4),
- we show that the manufacturing carbon footprint is a non-negligible factor in footprint reduction evaluations, and that – under certain conditions – minor footprint savings are possible when deploying *additional* sites where jobs are distributed according to the FTSFTW approach (Section 5),
- we show that larger relative footprint savings are possible when applying the FTSFTW scenario to distributed data centers where the nominal load is well below the maximum capacity (Section 6).

It should be noted that the theoretical model we present in Section 3 can be applied, with or without slight modifications, using other metrics than carbon footprint.

2. Related work

Next to the work already pointed out in the previous section, below are some earlier references and publication related specifically to the FTSFTW approach.

One of the first papers to suggest locating data centers near renewable energy sources is [14]. The primary reason given is that it is cheaper to transmit data over large distances than to transmit power. The paper does not discuss or explore this issue in any more detail.

The first paper to our knowledge to discuss and mathematically evaluate load distribution across data centers taking into account their energy consumption, energy cost (based on hourly electricity prices) and so-called low-footprint 'green energy' and high-footprint 'brown energy' is [10]. It presents and evaluates a framework for optimization-based request distribution, which is solved using heuristic techniques such as simulated annealing. The paper shows that it is possible to exploit green energy to achieve significant reductions in brown energy consumption for small increases in cost. It does not consider the manufacturing carbon footprint.

Similarly, in [15] load distribution across data centers is discussed, but only to optimize energy costs by exploiting energy price differences across regions.

In [5] the FTSFTW scenario is discussed in more detail and an Infrastructure as a Service (IaaS) approach is suggest to turn this in a viable business model. It outlines the main arguments for employing such a scenario. The key idea put forward is that the FTSFTW scenario provides a 'zero-carbon' infrastructure for ICT, thereby somewhat optimistically ignoring the potential contribution of the manufacturing carbon footprint.

The GreenStar Network project [16] is a proof-of-concept testbed for the FTSFTW strategy. The project started in 2010 and is deployed across the Canadian-based CANARIE research network and international partners distributed across the world. It consists of a number of small-scale 'nodes' powered by renewable energy (especially hydro, solar and wind power) which provide energy for the routers, switches and servers located at the node. Applications are running inside virtual machines, with multiple virtual machines per server, and are migrated live from node to node. The expected outcome of the project is a number of tools, protocols and techniques for deploying 'green' ICT services.

A framework for discovering carbon-minimizing resources in networks similar to those deployed by the GreenStar Network project, is described in [17], but again the manufacturing carbon footprint is not considered.