

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, tel: +32(0)9 33 14 977, fax: +32(0)9 33 14 899

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.

Keywords:

Green ICT, energy-efficiency, power consumption, distributed computing, grid computing, carbon footprint

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 to 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 disposition¹ for the U.S.A 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.

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

*Corresponding author

Email address: ward.vanheddeghem@intec.ugent.be (Ward Van Heddeghem)




Figure 1: Distributed data center

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][12][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].

Figure 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 CO₂-eq, unless otherwise indicated. CO₂-eq indicates CO₂-equivalent emissions, which is the amount of CO₂ 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 car-

bon 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 provided 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 suggested 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.

3. Theoretical Model

In this section we will outline the details of the scenario that we consider and develop a theoretical model for estimating its total carbon footprint. The quantification of the various parameters in our formulation will be done in Section 4.

To introduce our theoretical model, we consider the distributed simplified data center infrastructure that is shown in Figure 2. It consists of m equally-sized sites. Of these m sites, on average n sites are active. When a specific site becomes non-active, data and processing is moved to another active site, keeping the number of active data centers equal to n at all times.

At this point it is important to point out that, although we use the term data center, our model will be independent of the size of the data center. A data center site could be an energy-optimized building housing thousands of servers, or it could be as small as a single server. In the context of this paper, it might be helpful to think of a data center site as a computing node of any possible size.

Each site is powered by either LF or HF energy. The average availability of LF energy versus HF energy is considered equal, but uncorrelated, for each site. This availability ratio p might be the result of an average temporal availability of a specific renewable energy source (for example, solar or wind power), or specific service level agreements between the data center operator and the utility provider.

To reduce the total footprint, the usage of LF energy will be maximized by migrating operation of a data center powered by HF energy to a data center where LF energy is available. When no LF energy is available, HF energy will be used to guarantee service delivery.

The total carbon footprint F of the above described distributed data center infrastructure, averaged over a long-enough period, will be the sum of the manufacturing footprint F_m , the usage




Figure 2: Distributed data center infrastructure overview, consisting of $m=5$ sites with $n=3$ sites active. The independent LF energy availability per site is $p=0.6$

footprint F_u and the communication footprint F_c :

$$F = F_m + F_u + F_c \quad (1)$$

The manufacturing footprint will be the carbon emitted during the manufacturing of the sites and the equipment (servers, network equipment etc.) inside. The usage footprint will be the result of the electrical energy used during the use phase. The communication footprint will be the carbon emitted by migrating data and jobs from site to site. All three footprints will be expressed in g CO₂-eq.

Before we elaborate on each of these footprints, it is useful to point out the following assumptions we will make for our theoretical model:

- We assume each site in the distributed data center to be of uniform size.
- We assume instant site migration. That is, we assume that a migration takes no time and produces no extra overhead not accounted for in the communication footprint. If the migration frequency is relatively low (say, limited to a few times a day), this assumption will hold.
- We do not consider a surplus of LF energy. That is, if for example 4 out of 5 sites have LF renewable energy available, but we only require 3 sites for daily operation, the

electricity generated in the 4th site is 'wasted'. There is potential for using this energy for other less-critical purposes, or for selling or trading it for carbon credits. However, for simplicity and generality, our model does not take using surplus available power into account.

- We assume that a non-active data center site consumes no energy. While this is an optimistic assumption for large data centers, this is certainly feasible for micro-scale data centers consisting of a few servers (remember that, although we use the term data center, our model is independent of the data center size). The energy for a non-active site could be reduced to (nearly) zero by for example suspending all servers.

3.1. Usage Footprint

Let's call p the chance that a site is powered by LF energy. Let's call k the total number of data center sites that are powered by LF energy and P_k the chance of this number being k . This chance is given by the probability mass function of the binomial distribution:

$$P_k = \binom{m}{k} p^k (1-p)^{m-k} \quad (2)$$

Equation 2 can be understood intuitively as follows. The chance for exactly k sites powered by LF energy is p^k . The chance for the $m-k$ remaining sites to be *not* powered by LF energy is $(1-p)^{m-k}$. The number of ways to choose k sites out of a total of m sites is given by the binomial coefficient $\binom{m}{k}$ and can be calculated as $\binom{m}{k} = \frac{m!}{k!(m-k)!}$.

Given L the carbon footprint of the total usage phase of a single site when powered exclusively by LF energy and H the carbon footprint when powered exclusively by HF energy. The total usage footprint F_u for all sites is then:

If $k \geq n$ (that is, if LF energy is available in enough or more sites than required):

$$F_u = nL \quad (3)$$

Else:

$$F_u = (n-k)H + kL \quad (4)$$

Thus, using the chances of k being a certain value, the total usage footprint F_u becomes:

$$F_u = \sum_{k=n}^m [P_k nL] + \sum_{k=0}^{n-1} [P_k ((n-k)H + kL)] \quad (5)$$

The first term describes the weighted footprint if enough sites are powered by LF energy, the second term when this is not the

case. When substituting Equation 2 in 5 we get for the total usage footprint F_u of the distributed data center infrastructure:

$$F_u = nL \sum_{k=n}^m \left[\binom{m}{k} p^k (1-p)^{m-k} \right] + \sum_{k=0}^{n-1} \left[\binom{m}{k} p^k (1-p)^{m-k} ((n-k)H + kL) \right] \quad (6)$$

The usage footprint results exclusively from electrical energy. The emission intensity of electricity describes the GHG emissions in gram CO₂-eq per kWh. We use I_L and I_H to denote the emission intensity for LF and HF electricity respectively. With E_u the energy used by a single site during the entire use phase, L and H can thus be expressed as:

$$\begin{cases} L = I_L E_u \\ H = I_H E_u \end{cases} \quad (7)$$

3.2. Manufacturing Footprint

The total manufacturing footprint F_m is a function of the carbon footprint cost M for manufacturing one data center site, and the number of data centers sites m :

$$F_m = mM \quad (8)$$

As we will see in Section 3.4, it is convenient to consider the manufacturing fraction f , which is the ratio of the manufacturing carbon footprint M of a single site over the usage carbon footprint H of a single site:

$$f = \frac{M}{H} \quad (9)$$

Equipment where the manufacturing emits less GHG than the typical GHG emitted during its use phase will have a manufacturing fraction $f < 1$.

Given Equation 9, we can rewrite Equation 8 as:

$$\begin{aligned} F_m &= mfH \\ &= mfI_H E_u \end{aligned} \quad (10)$$

Note that we considered the equipment to be manufactured with HF energy, by expressing M as a function of H instead of L .

Reference	Description	Manufacturing phase	Use phase (4 years)	f
PE International [18]	Simple office server ^a	500 kg CO ₂ -eq/unit	1030 kg CO ₂ -eq	0.49
Malmodin ITU [19]	PC ^a	400 kg CO ₂ -eq/unit	640 kg CO ₂ -eq	0.63
Malmodin ITU [19]	Server	500 kg CO ₂ -eq/unit	5200 kg CO ₂ -eq	0.10
Malmodin, Moberg [20]	Data centers ^b	10 Mton CO ₂ -eq in 2007	108 Mton CO ₂ -eq in 2007	0.09

^a Overhead power in use phase not included (PUE=1). See text for more information

^b This includes data center equipment and buildings. Data based on 10 million new servers and 35 million servers in use; this translates roughly to a use phase of 4 years. Use phase emission intensity in [20] = 0.6 kg CO₂-eq/kWh

Table 1: Manufacturing fraction values according to different studies

3.3. Communication Footprint

Migrating jobs or data across data centers incurs an extra amount of carbon emissions. This will mainly be due to the energy consumed for (a) the transportation over an optical network, (b) the preparation and duration of the migration and (c) switching the data center to the non-active state or vice versa. In this section we show that the overhead of the above three factors is negligible with respect to the carbon emitted in the manufacturing and use phase, and can thus be ignored for now.

Data centers are typically connected by optical networks. Power consumption in the optical core network is dominated by the IP router power consumption, with high-end IP routers consuming in the order of 10 W/Gbps [21]. Accounting for redundancy, cooling and power supply overhead, and client and network interface, we have approximately 100 W/Gbps, or an energy of $2.7 \cdot 10^{-5}$ kWh needed to transport one Gbit.

Further, we assume two migrations per site once a day, i.e. one inbound migration and one outbound migration. We consider each server in a data center site to be capable of running four virtual machines, with each virtual machine to be about 10 Gbyte in size. For each server's data to be migrated, this totals to 640 Gbit/day. Considering a server use phase of 4 years, this sums up to 934 000 Gbit per use phase. Using our estimation from above, this requires approximately 26 kWh of energy. With a world-average emission intensity of 500 g CO₂-eq/kWh, this results in about 13 kg CO₂-eq emitted due to migration (for one server, during its entire use phase). This equals to less than 3% of the current manufacturing footprint of a server (about 500 kg CO₂-eq, see Table 1), or about 0.5% of the current total carbon emissions.

With respect to the energy overhead induced by migration preparation and duration, transmitting our exemplary 640 Gbit/day would take less than 15 minutes per day over a 1 Gbps link. This accounts for only about 1% of the time.

Likewise, as the daily migration frequency is low, the time and energy overhead to switch a data center from the active to non-active state (or vice versa) should be relatively low as well. Also, the active/non-active switchover time will probably depend on the kind of jobs and data that the data center is running.

Although the above estimate is based on the current situation of the average absolute carbon footprint of servers and current virtualization technology, we feel that it is a fair assumption for

current and short term future to neglect the contribution of the communication footprint F_c to the total footprint.

3.4. Total Footprint

Combining Equation 6 and 10, the total footprint is given by:

$$\begin{aligned}
F &= mfI_H E_u \\
&+ nL \sum_{k=n}^m \left[\binom{m}{k} p^k (1-p)^{m-k} \right] \\
&+ \sum_{k=0}^{n-1} \left[\binom{m}{k} p^k (1-p)^{m-k} ((n-k)H + kL) \right] \quad (11)
\end{aligned}$$

The above equation depends on the value of E_u , the single site usage energy. This value will vary depending on the data center size and type, and on the jobs and data processed. We can eliminate this parameter, if we normalize the total footprint over the single site usage energy E_u .

By doing so, we can conveniently express this total normalized footprint F_{norm} as a function of the LF energy emission intensity I_L , the HF energy emission intensity I_H and the fraction f :

$$\begin{aligned}
F_{norm} &= \frac{F}{E_u} \\
&= mfI_H \\
&+ nI_L \sum_{k=n}^m \left[\binom{m}{k} p^k (1-p)^{m-k} \right] \\
&+ \sum_{k=0}^{n-1} \left[\binom{m}{k} p^k (1-p)^{m-k} ((n-k)I_H + kI_L) \right] \quad (12)
\end{aligned}$$

We now have a metric for the carbon footprint which is independent from the data center size and type, and with unit [g CO₂-eq/kWh].

4. Parameter Quantification

Our model constructed in the section above consists of a number of parameters. In this section we discuss realistic values for each of these parameters.

4.1. Manufacturing fraction (f)

The manufacturing fraction represents the ratio between the manufacturing carbon footprint and the usage carbon footprint. Detailed life cycle analysis (LCA) studies that report on the carbon emissions of data centers during the manufacturing phase and the use phase are scarce. Moreover, the resulting manufacturing fraction is influenced by the use phase lifetime of the equipment and the emission intensity of the energy used during the use phase. In addition, it is important to know if reported use phase values include power consumed for overhead such as cooling. This overhead is typically expressed by the power usage effectiveness (PUE). For example, a PUE of 2 (a typical accepted value for data centers²) indicates that for each Watt consumed by useful equipment such as servers and switches an additional Watt is consumed through overhead.

Table 1 lists emission values and the derived manufacturing fraction f according to a number of studies. All data, except for the 'Simple office server' and the 'PC', includes overhead power consumption. For the 'Simple office server' probably no overhead is included ([18] isn't completely clear on this); correcting for this with a PUE of 2, the use phase power consumption doubles and thus the manufacturing fraction value halves, bringing the values roughly in line with the other data.

Based on the data in Table 1 we will use, unless otherwise specified, a value of $f=0.25$.

4.2. High-footprint energy emission intensity (I_H)

The parameter I_H indicates the emission intensity of regular (HF) electrical energy. As already stated, the emission intensity indicates the amount of GHGs emitted for each kWh of electrical energy, and is typically expressed in grams of CO₂-eq per kWh.

The value for I_H differs from country to country, and for larger countries even from region to region, depending on the primary energy sources (such as coal or gas) and technologies (such as open cycle gas turbines or combined cycle gas turbines) used for generating electricity, see for example Table 2³.

For this paper, we will consider the world average value of 500 g CO₂-eq/kWh.

²Recently deployed high-capacity data centers with a focus on energy efficiency show much lower PUE values, such as Google claiming to reach a yearly average of 1.16 at the end of 2010 [22]. However, as the LCA data is based on 2007 estimates, the for that year typically accepted PUE value of 2 is used [23].

³The table reports the CO₂ emissions instead of the CO₂-eq emission (which takes a number of other GHGs into account). However differences are minor and irrelevant for our study

Region	Intensity [g CO ₂ /kWh]
World	502
United States	535
Canada	181
European Union	351
China	745
India	968

Table 2: Average CO₂ emissions per kWh from electricity and heat generation for a number of countries and regions, data for 2008 [24]

4.3. Low-footprint energy emission intensity (I_L)

The emission intensity I_L for low-footprint electricity is obviously lower than the regular HF energy emission intensity I_H . Indicative, Figure 3 lists the estimated emission intensity for a number of low-footprint sources (typically renewable energy such as hydro, wind or solar power), as reported by [9]. Roughly similar numbers are given in the slightly older study of [25].

In this paper, we assume a state-of-the-art LF energy emission intensity of 10 g CO₂-eq/kWh.




Figure 3: Lower and upper emission intensity estimates for various low-footprint sources [9] (CSP: Concentrated Solar Power)

4.4. Low-footprint energy availability (p)

The parameter p represents the chance of each site being powered by LF energy. For example, with $p=0.6$, each site has an independent chance of 60% to be powered by LF energy at any point in time. Or otherwise put, 60% of the time, each site will be powered by LF energy.

While it might seem tempting to try to relate the value for p to the availability of a specific LF energy source (say, wind energy), this is not necessary for our model. After all, the availability of LF energy sufficient for powering a data center site will largely be a matter of monetary cost. This cost will be reflected either in the negotiated service level agreement (SLA) with the utility provider, or in the cost to install the required capacity of LF energy sources to deliver the required nominal

power even during periods of low availability of e.g. sun or wind. Thus, a higher value for p will usually require higher investments. Note that it is key for the validity of our footprint model to know what kind of power (LF or HF) is used at what point in time, so as to be able to migrate the data to a different site if needed (and if possible).

We assume $p=0.6$, as we will see later that this results in maximum savings.

5. Case Study I: The Added Distributed Data Centers (ADD) Scenario

Can we reduce the footprint of a regular data center, by distributing additional sites across the globe as to benefit from uncorrelated and potentially complementary availability of renewable energy sources which offer a lower usage footprint? This is the question we will examine in this section. We refer to this scenario as the *Added Distributed Data centers (ADD) scenario*.

Consider a data center that requires $n=3$ sites for daily operation. Each site has an LF energy availability of $p=0.6$, and we consider the current estimation for the manufacturing fraction $f=0.25$. Since we want to reduce the footprint of the complete data center, we would like to be able to run our applications on three data centers that have LF energy available. The chance of success increases with an increased number of data centers to choose from, that is, if we increase the total number of sites m to a value higher than 3.

Figure 4 shows the use phase, manufacturing phase and total footprint as we increase the total number of data centers m beyond 3. With each additional data center, the use phase footprint decreases as a result of the increased chance of finding a data center that runs on LF energy. Initially, this decrease is large enough to make up for the linearly increasing manufacturing footprint, resulting in a decreasing total footprint. However, when the number of data centers is approximately the double of the number of data centers required, the total footprint increases and eventually overtakes the first scenario footprint.

Taking the first scenario (where $m=n=3$) as a baseline, we see initial footprint savings until too much data centers are deployed, resulting in a net loss. Taking the first scenario as the baseline makes sense, since this corresponds to the current practice of operating a number of sites with a mix according to p of LF and HF energy, *without* migrating data or processing capacity based on LF energy availability.

5.1. Influence of manufacturing fraction (f)

As we have seen in the above case, the usage footprint reduction was initially able to make up for the linearly increasing manufacturing footprint. What if the manufacturing fraction f is higher, say $f=0.5$? Figure 5 shows the normalized footprint (upper figure) and relative savings (lower figure) for different values of f .




Figure 4: The total normalized footprint F_{norm} and corresponding relative emission savings as a function of the total number of data centers m . Savings are calculated with respect to the baseline scenario. (Parameter values: $n=3$, $f=0.25$, $p=0.6$, $I_L=10$ g CO₂-eq/kWh, $I_H=500$ g CO₂-eq/kWh and $E_u=1$)

Clearly, footprint reduction becomes smaller and even impossible for higher values of f . Even more so, our current rough estimate of $f=0.25$ seems critical: with a slightly higher value for $f=0.3$ savings are almost negligible (a mere optimistic 5%) and might be completely annihilated if we take more subtle factors (such as the migration footprint and management overhead) into account.

In the inverse case, for lower values of f the savings increase. At the utopian case of having manufacturing for free ($f=M=0$), savings are obviously maximal and converge to the usage footprint cost nL .

5.2. Influence of low-footprint energy availability (p)

Perhaps counterintuitive, an increase of LF energy availability of p towards 100% does not unconditionally result in additional savings. While the footprint indeed decreases monotonic with an increase of p (because the usage footprint becomes smaller), the baseline scenario footprint (where $m=n$) will also decrease.

Figure 6 shows that for the scenario $n=3$, $m=6$ (i.e., twice as much data centers as required for daily operation) the savings are maximum around $p=0.5$ to 0.6 . For $p=0$ there is a net loss due to the increased manufacturing footprint not yet being offset by a greener usage footprint. For $p=1$ the baseline scenario runs entirely on LF energy whereas the FTSFTW approach has an increased manufacturing footprint due to the extra sites deployed.

As we have already argued that p will be cost driven, a case-based cost study will have to find the optimal value for p . In retrospect, this also explains our decision for taking $p=0.6$.

5.3. Influence of n and m values

Because of the binomial coefficient, we cannot simply generalize the footprint savings obtained for e.g. $n=3$ and $m=6$ to


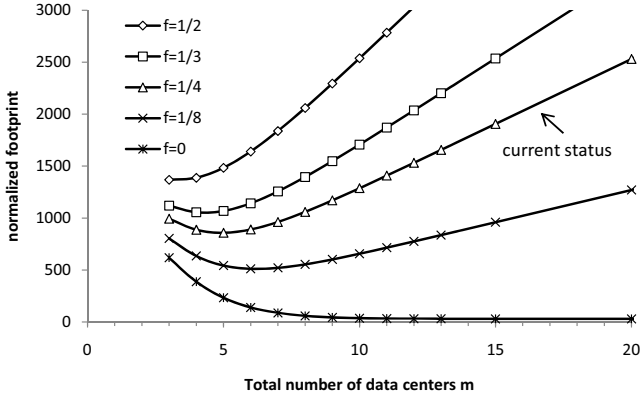


Figure 5: The total normalized footprint F_{norm} and relative emission savings for $n=3$ as a function of m for different manufacturing fractions f




Figure 6: The total normalized footprint and the relative savings (with respect to the baseline scenario where $m=n=3$) as a function of the LF energy availability p ($n=3$, $m=6$ and $f=0.25$)

apply to any other combination of n and m with the same ratio, e.g. $n=1$ and $m=2$, or $n=10$ and $m=20$.

For higher values of n , footprint savings already occur for higher (i.e., worse) manufacturing fractions. For example, when we consider $n=10$ (see Figure 7), already for $f=0.5$ minor savings are available (2% maximum), whereas for the previous case where $n=3$ this was not the case (see Figure 5). Because of the higher number of sites, the chance for finding enough sites where LF energy is available has increased. It should be noted that the *total* footprint will have increased as well.

This finding suggests to favor a large number of small, distributed data center sites, over a few large ones. However, in that case, care should be taken that the combined manufacturing footprint of the small sites is not larger than the manufacturing footprint of the few larger sites. Taking the idea to extremes, large-scale distributed computing projects such as Folding@Home [26] where small consumer entertainment devices are involved [27] might be a perfect fit, if both the manufacturing footprint and usage footprint (standby power consumption issues etc.) from these devices is small enough.




Figure 7: Relative emission savings for $n=10$ (all other parameters are equal as before)

5.4. Influence of emission intensity difference

The HF emission intensity (500 g CO₂-eq/kWh) en LF emission intensity (10 g CO₂-eq/kWh) that we consider in this paper following our findings in Section 4.2 and 4.3 are relatively large in difference; I_L is only 2% of the I_H . In some countries or regions, the regular emission intensity is substantially lower (or higher) than the world average value, as can be seen in Table 2. Will the FTSFTW approach still be sustainable under those conditions?

Figure 8 shows the relative savings with changing values of I_H . It is immediately clear from this figure that for values below the world average, the savings quickly become negligible. For emission intensities below the average European value, savings become negative, i.e. *more* carbon dioxide will be emitted. On the contrary, for geographical regions where the regular elec-

tricity has high emission intensities (such as China and India), the savings offered by FTSFTW are much higher.

Note that we consider the manufacturing carbon footprint cost M (see equation 8) to be fixed, even with changing I_H value. This means that in this case we have fixed the instance of I_H in equation 10 to the world-average emission intensity. Fixing the manufacturing footprint makes sense, as it represents the case where the equipment remains manufactured as before, but is used in a region with a different HF energy emission intensity.

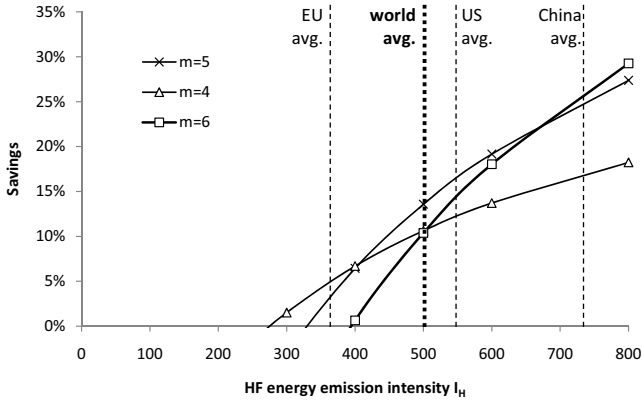


Figure 8: Relative emission savings as a function of the HF energy emission intensity I_H . (for $f=0.25$, $p=0.6$ and $I_L=10$ g $\text{CO}_2\text{-eq/kWh}$)

Similarly, we can also consider different values for the LF energy emission intensity I_L . The value of $I_L=10$ g $\text{CO}_2\text{-eq/kWh}$ we assumed in Section 4.3, is based on state-of-the art renewable energy, typically from wind turbines. For other energy sources with higher emission intensities, the savings will obviously be smaller.

Figure 9 shows the savings for increasing values of I_L , with the HF energy emission intensity fixed at 500 g $\text{CO}_2\text{-eq/kWh}$. As can be seen, the savings rapidly dwindle, to the point where they become marginal. As such, using less emission-saving renewable energy sources such as solar PV installation should be evaluated carefully if the main goal is saving on total carbon emission by employing the ADD scenario.

To summarize, with current estimates for the manufacturing and usage footprint, carbon emission savings up to around 14% are possible by deploying additional data center sites. Actual savings depend mainly on the manufacturing fraction (lower is better), the LF energy availability (optimum around 50–70%) and the number of sites deployed (optimum around 1.5 to 2.5 times as much data centers as required for daily operation). For geographical regions with higher HF emission intensities, the possible savings by employing the ADD scenario are much higher than 14%; likewise, for intensities below the world average savings quickly turn negative.




Figure 9: Relative emission savings as a function of the LF energy emission intensity I_L . (for $f=0.25$, $p=0.6$ and $I_H=500$ g $\text{CO}_2\text{-eq/kWh}$)

6. Case Study II: The Low Load Redistribution (LLR) Scenario


The main conclusion from the above scenario is that the manufacturing carbon footprint is a non-negligible factor, and should be taken into account when evaluating potential carbon footprint savings. However, there are cases where the manufacturing footprint is already expended. Data centers are not constantly running at peak capacity, but instead operate at a nominal load well below the peak capacity, typically servers operate most of the time between 10 and 50 percent of their maximum utilization levels [28]. We could redistribute the load using the FTSFTW approach, resulting in what we will refer to as the *Low Load Redistribution (LLR) scenario*.

Regular approach — Figure 10-a shows the regular approach (without applying LLR). The load is equally distributed among the different sites. To calculate the total carbon footprint, we consider a data center with peak capacity m to run at nominal load n . We assume unused servers to be powered down. The total footprint of a data center running at this nominal load is then:

$$F_{nominal} = F_u + F_m = n(pL + (1-p)H) + mM \quad (13)$$

LLR approach — What would happen if we apply the FTSFTW approach to optimally distribute processing to sites where LF energy is available (Figure 10-b)? We can use Equation 11 or 12 to calculate the footprint in that case as well, with m representing the peak capacity, and n representing the (varying) nominal load.

Figure 11 plots the footprint for both scenarios for a distributed data center consisting of 5 sites ($m=5$), for an increasing load (i.e., n increasing from 0 to m). The LF energy availability p per site has been taken equal to 0.6. As can be seen, for the nominal load being half of the peak capacity, savings around 20% are possible by employing FTSFTW. These are savings



(a) Regular scenario: 3/5 load distribution under the regular approach (b) LLR scenario: 3/5 load distribution under the FTSFTW approach

Figure 10: Nominal load distribution in a distributed data center. (a) shows the regular scenario where a nominal load of 60% is distributed equally over all data center sites. (b) shows the LLR scenario, where the same nominal load is distributed according to the FTSFTW approach, resulting in an optimal usage of sites with LF energy availability




Figure 11: Relative footprint (with respect to the maximum load) of a distributed data center running at various loads both under a regular scenario and a LLR scenario ($m=5, p=0.6$). The 'Savings, total' are the relative savings over the total footprint (both manufacturing F_m and usage F_u). The 'Savings, usage' are the relative savings over the usage footprint only.

over the total footprint, that is, the sum of the use phase and manufacturing phase footprint.

If we only consider the savings over the usage phase, which would be an equally valid approach since the manufacturing phase has no savings, the savings are as high as 90% when running at 20% of the capacity and still reach more than 60% when running at half the peak capacity.

The savings itself vary for different values of p . This is shown in Figure 12.




Figure 12: Savings for various values of p

It is important to remark that from the above results we should not conclude to design distributed data centers to run well below their maximum capacity. This results in a large total manufacturing carbon footprint. First, and foremost, data center capacity should be scaled to their nominal loads as much as possible, taking into account such factors as redundancy and peak loads. Once this is done, carbon emissions can be reduced using the LLR scenario outlined above.

7. Conclusions

The carbon footprint from data centers is significant, and growing. Besides improvements in energy-efficiency, the use of low footprint energy (typically from renewable energy sources such as wind or solar power) is key to reducing data center carbon

emissions. Data centers are in a unique position to overcome some of the issues currently associated with renewable energy sources. They can be located near renewable energy sites, and jobs and data can be migrated from site to site as renewable energy — intermittent by nature — comes and goes. This approach has been referred to as follow the sun/follow the wind.

In this paper, we researched if carbon emissions can be reduced by applying this technique to take advantage of the resulting increased availability of low footprint renewable energy. To this purpose, we have build a mathematical model to calculate the carbon footprint of such a distributed data center infrastructure that is powered by a mix of low-footprint (LF) and high-footprint (HF) energy. We have shown that for footprint reduction the manufacturing carbon footprint of data centers is a critical parameter to consider. Based on the available LCA data for data centers, footprint savings in the order of 14% over the total footprint are possible by deploying additional data center sites to take advantage of the resulting increased available of LF energy. Reductions of the manufacturing footprint relative to the usage footprint will lead to improved savings. However, a number of factors heavily influence the actual savings, which could easily turn into an *increased* carbon footprint if not evaluated carefully. As the savings are strongly influenced by the HF electrical emission intensity, it is of no use to deploy the follow the sun/follow the wind approach in regions with emission intensities below the current world average value. And, consequently, it makes more sense to use the approach in regions with high emission intensities. Carbon footprint savings also depend on the LF energy availability per site: optimal availability varies for different configurations, but is in the order of 50–70%. Optimum savings can be gained at architectures that deploy around 1.5 to 2.5 times as much data centers as required for daily operation.

Bigger savings — up to 60–90% — are possible by applying the follow the sun/follow the wind strategy to data centers where the nominal load is well-below the peak capacity.

Finally, it should be noted that our model is not restricted to carbon footprint metrics. It can easily be used or modified to evaluate other metrics. For example, the low and high emission intensities can be replaced by low and high energy prices (requiring an appropriate quantification of the manufacturing fraction in that case) to evaluate the cost benefits in the light of fluctuating energy prices. However, this is outside the scope of this paper.

8. Acknowledgments

The authors would like to thank Koen Casier, Kevin Mets and Tom Verschueren for their helpful advice and discussions. We also thank the Greenstar Network partners for their valuable comments. The work described in this paper was carried out with the support of the BONE project (Building the Future Optical Network in Europe) and the TREND project (Towards Real Energy-efficient Network Design), both a Network

of Excellence funded by the European Communitys Seventh Framework; the IBBT-project GreenICT and the STRONGEST project funded by the European Communitys Seventh Framework Programme FP7/2007-2013 under grant agreement n 247674.

References

- [1] M. Pickavet, W. Vereecken, S. Demeyer, P. Audenaert, B. Vermeulen, C. Develder, D. Colle, B. Dhoedt, P. Demeester, Worldwide energy needs for ICT: the rise of power-aware networking, in: 2008 2nd International Symposium on Advanced Networks and Telecommunication Systems, p. 3 pp.
- [2] EIA, International Energy Statistics, total electricity net consumption, <http://tonto.eia.doe.gov/cfapps/ipdbproject/IEDIndex3.cfm?tid=2&pid=2&aid=2>, 2010.
- [3] J. Koomey, Estimating total power consumption by servers in the US and the world, 2007.
- [4] L. Deboosere, P. Simoens, J. De Wachter, B. Vankeirsbilck, F. De Turck, B. Dhoedt, P. Demeester, Grid design for mobile thin client computing, Future Generation Computer Systems (2010).
- [5] S. Figuerola, M. Lemay, V. Reijts, M. Savoie, B. St Arnaud, Converged Optical Network Infrastructures in Support of Future Internet and Grid Services Using IaaS to Reduce GHG Emissions, Lightwave Technology, Journal of 27 (2009) 1941–1946.
- [6] A. Ipakchi, F. Albuyeh, Grid of the future, Power and Energy Magazine, IEEE 7 (2009) 52–62.
- [7] T. Ackermann, G. Andersson, L. Söder, Distributed generation: a definition, Electric Power Systems Research 57 (2001) 195–204.
- [8] EIA, Total annual losses estimates related to electrical transmission and distribution, [http://tonto.eia.doe.gov/ask/electricity/_faqs.asp#electric\\$rates2](http://tonto.eia.doe.gov/ask/electricity/_faqs.asp#electric$rates2), 2007.
- [9] M. Jacobson, Review of solutions to global warming, air pollution, and energy security, Energy & Environmental Science 2 (2009) 148–173.
- [10] K. Le, R. Bianchini, M. Martonosi, T. Nguyen, Cost-and energy-aware load distribution across data centers, Proceedings of HotPower (2009).
- [11] C. Clark, K. Fraser, S. Hand, J. G. Hansen, E. Jul, C. Limpach, I. Pratt, A. Warfield, Live migration of virtual machines, in: NSDI'05: Proceedings of the 2nd conference on Symposium on Networked Systems Design & Implementation, USENIX Association, Berkeley, CA, USA, 2005, pp. 273–286.
- [12] F. Travostino, P. Daspit, L. Gommans, C. Jog, C. De Laat, J. Mambretti, I. Monga, B. Van Oudenaarde, S. Raghunath, Seamless live migration of virtual machines over the MAN/WAN, Future Generation Computer Systems 22 (2006) 901–907.
- [13] G. Kecskemeti, G. Terstyanszky, P. Kacsuk, Z. Neméth, An approach for virtual appliance distribution for service deployment, Future Generation Computer Systems 27 (2011) 280–289.
- [14] A. Hopper, A. Rice, Computing for the Future of the Planet, Phil. Trans. R. Soc. A (2008).
- [15] A. Qureshi, R. Weber, H. Balakrishnan, J. Guttag, B. Maggs, Cutting the electric bill for internet-scale systems, in: Proceedings of the ACM SIGCOMM 2009 conference on Data communication, ACM, pp. 123–134.
- [16] The Greenstar Network Project, Project website: <http://www.greenstarnetwork.com/>, 2010.
- [17] A. Daouadji, K.-K. Nguyen, m. Lemay, M. Cheriet, Ontology-based Resource Description and Discovery Framework For Low Carbon Grid Networks, in: SmartGridComm 2010.
- [18] C. Hermann, Environmental footprint of ICT equipment in manufacture, use and end of life, in: ECOC 2008, Brussels, p. 24.
- [19] J. Malmmodin, Global Carbon Footprint of ICT: Method and Results, ITU FG ICT & CC, Geneva, 2008.
- [20] J. Malmmodin, A. Moberg, D. Lundn, G. Finnveden, N. Lvehagen, Greenhouse Gas Emissions and Operational Electricity Use in the ICT and Entertainment & Media Sectors, Journal of Industrial Ecology (2010).
- [21] W. Van Heddeghem, M. De Groote, W. Vereecken, D. Colle, M. Pickavet, P. Demeester, Energy-efficiency in telecommunications networks : link-

- by-link versus end-to-end grooming, in: Conference on Optical Network Design and Modeling, 14th, Proceedings, IEEE, 2010, p. 6.
- [22] Google, Data center efficiency measurements, <http://www.google.com/corporate/datacenter/efficiency-measurements.html>, 2011.
 - [23] U.S. Environmental Protection Agency, Report to Congress on Server and Data Center Energy Efficiency, http://www.energystar.gov/ia/partners/prod_development/downloads/EPA_Datacenter_Report_Congress_Final1.pdf, 2007.
 - [24] IEA, CO2 Emissions from Fuel Combustion - highlights, <http://www.iea.org/co2highlights/>, 2010.
 - [25] M. Pehnt, Dynamic life cycle assessment (LCA) of renewable energy technologies, *Renewable Energy* 31 (2006) 55–71.
 - [26] Stanford University, Folding@home, Project page: <http://folding.stanford.edu/>, 2000.
 - [27] E. Luttmann, D. Ensign, V. Vaidyanathan, M. Houston, N. Rimon, J. Øland, G. Jayachandran, M. Friedrichs, V. Pande, Accelerating molecular dynamic simulation on the cell processor and Playstation 3, *Journal of computational chemistry* 30 (2009) 268–274.
 - [28] L. Barroso, U. Holzle, The case for energy-proportional computing, *Computer* 40 (2007) 33–37.