On Satisfying Green SLAs in Distributed Clouds

Ahmed Amokrane*, Mohamed Faten Zhani†, Qi Zhang†, Rami Langar*, Raouf Boutaba†, Guy Pujolle*

*LIP6 / UPMC - University of Pierre and Marie Curie; 4 Place Jussieu, 75005 Paris, France
† University of Waterloo; 200 University Ave. W., Waterloo, ON, Canada
Email: ahmed.amokrane@lip6.fr; mfzhani@uwaterloo.ca; rami.langar@lip6.fr; rboutaba@uwaterloo.ca; guy.pujolle@lip6.fr

***Abstract*—With the massive adoption of cloud-based services, high energy consumption and carbon footprint of cloud infrastructures have become a major concern in IT industry. Consequently, many governments and IT advisory organizations have urged IT stakeholders (i.e., cloud provider and cloud customers) to embrace green IT and regularly monitor and report their carbon emissions and put in place efficient strategies and techniques to control the environmental impact of their infrastructures and/or applications.**

Motivated by this growing trend, we investigate, in this paper, how cloud providers can meet Service Level Agreements (SLAs) with green requirements. In such SLAs, a cloud customer requires from cloud providers that carbon emissions generated by the leased resources should not exceed a fixed bound. We hence propose a resource management framework allowing cloud providers to provision resources in the form of Virtual Data Centers (VDCs) (i.e., a set of virtual machines and virtual links with guaranteed bandwidth) across a geo-distributed infrastructure with the aim of reducing operational costs and green SLA violation penalties. Extensive simulations show that the proposed solution maximizes the cloud provider's profit and minimizes the violation of green SLAs.

I. Introduction

With the rapid development of cloud computing technologies, data centers have become a popular platform for delivering large-scale online services such as content delivery, social networking and e-commerce. However, the rapid expansion of cloud infrastructures in recent years have also raised serious concerns regarding their energy consumption and environmental impact. Recent reports [1] have revealed that the Information and Communication Technologies (ICT) account for 3% of the world's carbon emissions. Data centers by themselves accounts for about 10% of the ICT emissions worldwide.

Motivated by these observations, the ICT sector is witnessing an upward move towards greening cloud infrastructures and services driven by several governmental regulations and marketing considerations. For instance, a recent study [2] showed that the firms' value would decrease significantly if it has high carbon footprint or even if it withholds information about its carbon emission rates. As a result, many IT companies are voluntarily disclosing their carbon emissions and regularly reporting their efforts towards deploying environmental-friendly solutions and services [3]. At the same time, governments are imposing taxes on carbon emissions in the hopes of pushing further this shift towards the adoption of green sources of energy and the reduction of carbon footprint [4].

In current cloud environments, there are mainly two stakeholders: (1) *cloud providers* (CPs) that typically own distributed infrastructures and lease their resources in an on-demand manner to different *Service Providers* (SPs); (2) SPs use these resources to deploy their services and offer them to Internet end-users. Recent research proposals and cloud offerings [5] are advocating to offer these resources in the form of Virtual Data Centers (VDCs), i.e., a set of VMs and virtual links with guaranteed bandwidth.

Typically, CPs are responsible for allocating resources for VDCs across their distributed clouds with the goal of minimizing operational costs and maximizing the infrastructure environmental friendliness by increasing the usage of green energy [6]. However, recently, SPs were also required to take into account environmental objectives and ensure that their services are produced with the smallest carbon footprint. Many advisory boards and commissions (e.g., Open Data Center Alliance [7] and SLA Expert Subgroup of the Cloud Selected Industry Group of the European Commission [8]) are pushing towards defining green SLAs in which SPs require their CPs to limit the carbon emissions generated on their behalf. Recently, some research works advocated providing Green SLAs in the context of HPC clouds [9]–[13].

Typically, the green SLA terms require either to limit the carbon emissions generated by SPs services [9]–[12] or to set a minimum amount of renewable power to be consumed by the resources allocated to SPs [13]. However, these proposals do not consider the allocation of network resources (virtual links) and aim only to allocate VMs within a single data center.

In this paper, we investigate how a CP can meet an SLA with green requirements. In particular, we consider SLAs that specify a limit on the carbon emission generated by each service provider's VDC. We, hence, propose Greenslater, a holistic framework that orchestrates the provisioning and the resource optimization for the multiple VDCs deployed across a distributed infrastructure. From the CP's point of view, the objective is to maximize revenue while minimizing operational costs and the potential green SLA violation penalties. Greenslater takes advantage of the variability in space and time of the available renewables and electricity prices in different data centers to reduce

the carbon footprint and costs. It provisions VDCs and dynamically optimize resource allocation over time while fulfilling the green SLA terms. Through extensive simulations, we show that the proposed framework maximizes the CP's profit and also the usage of renewable power while minimizing SLA violation penalty.

The remainder of this paper is organized as follows. Section II surveys the related works. Section III defines green SLAs and presents the proposed management framework. The mathematical formulation of the VDC embedding problem across distributed infrastructures that considers green SLAs is then presented in Section IV. Section V gives a detailed description of the proposed algorithms for VDC admission control and dynamic resource allocation and optimization. Section VI discusses simulation setup and results. Finally, we conclude the paper in Section VII.

II. Related work

In the last few years, a large body of work has addressed the problem of reducing energy consumption and carbon footprint in cloud environments. In the following, we first survey the literature on green management in the cloud and then we focus on the proposals that advocated implementing green SLAs between cloud and service providers.

A. Green management in the cloud

Recently, several systems have been proposed to map VDCs onto a single data center with the goal of reducing energy consumption. For instance, Zhani et al. [14] proposed VDC Planner, a resource management framework that leverages dynamic VM migration to increase CP's revenue while minimizing energy consumption. Unfortunately, these solutions are designed to manage a single data center and hence do not consider the variability over time and between different locations of the electricity prices and the availability of green sources of energy.

A plethora of techniques have been also proposed to allocate resources across geographically distributed data centers in order to reduce energy costs [15]–[17], minimize the infrastructure's carbon footprint [18], [19] or achieve both objectives [6], [20], [21]. For instance, Xin et al. [22] proposed an algorithm that uses *minimum k-cut* to split a VDC request into partitions before assigning them to different locations so as to balance the load among different data centers. In [6], we proposed Greenhead, a framework for VDC embedding across distributed infrastructures that aims at maximizing cloud providers' revenue while cutting down the carbon footprint of the infrastructure. Unfortunately, the solutions above use static mapping and do not perform any dynamic resource optimization over time. They also do not consider green SLAs and hence do not guarantee any limit on carbon emissions of the resources leased by each SP.

B. Green SLA in the cloud

Green SLAs stipulate that SPs are able to require their cloud providers to guarantee that the leased resources are environmental friendly. In other words, SPs can explicitly specify green constraints like, for instance, an upper limit on carbon emissions produced by the resources they lease.

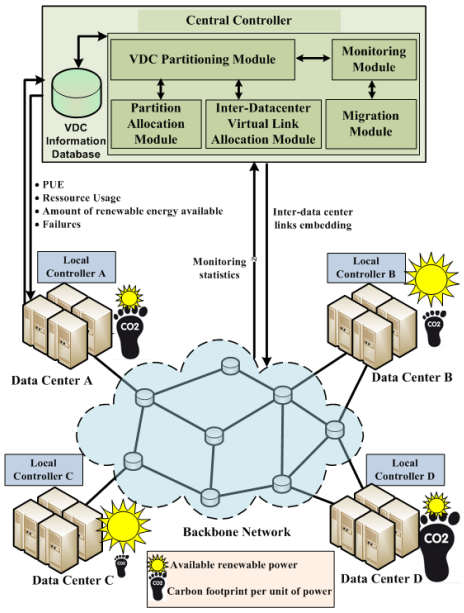

Fig. 1: Proposed framework

Providing green SLAs has been originally proposed back in 2010 by Laszewski et al. [9] and then quickly adopted and supported in several research works [11], [13], [23]. For instance, Haque et al. [13] considered an SLA that specifies the proportion of green power that the HPC provider should use to run the job (e.g., *x% of the job should run on green power*). The HPC provider has to pay a penalty to SPs if the green terms of the SLA are not satisfied. Similarly, Wang et al. [23] proposed an approach where SPs can define in the SLA constraints for their submitted tasks that limit the carbon emissions and consumed power. From the CP perspective, the goal of this framework is to schedule parallel tasks such that the green SLAs are satisfied.

However, the solutions above do not consider bandwidth requirements between VMs and are designed to manage resources within a single data center. Our work consider a more general scenario with multiple data centers and where the network requirements are explicitly specified in the VDC request.

III. System Architecture

In this section, we present the design architecture of the proposed solution and we discuss the definition of the Green SLA terms and how to enforce them in a distributed environments.

A. Architecture Overview

As shown in Fig. 1, we consider a distributed infrastructure consisting of multiple data centers located in different regions and interconnected through a backbone network. The entire infrastructure (including the backbone network) is assumed to be owned and managed by the same CP.

SPs send VDC request specifications to the CP, which has the responsibility of allocating the required resources.

Naturally, the CP will make use of its distributed infrastructure with the objective of maximizing its revenue and minimizing energy costs and carbon footprint; this is where our proposed management framework, Greenslater, comes into play. Greenslater is composed of two types of management entities: i) a Central Controller that manages the entire infrastructure and ii) a Local Controller deployed in each data center to manage the data center's internal resources (i.e., resource allocation for VMs and virtual links inside the data center).

The central controller consists of a number of components. The *Partitioning Module* is in charge of dividing a VDC request into partitions such that inter-partition bandwidth is minimized. The *Partition Allocation Module* is then responsible for running an admission control algorithm for every received VDC request, and assigns the partitions, in case of accepted requests, to data centers based on runtime statistics collected by the monitoring module and the estimation of available renewable power. The *Inter-data center Allocation Module* allocates resources for the virtual links spanning the backbone network. Finally, the *Migration Module* dynamically relocates VDC partitions in such a way to follow renewables and reduce the carbon footprint. The *Monitoring Module* monitors and collects information about the status of physical and virtual infrastructures and stores them into *VDC Information Base*.

B. Green SLA Definition

As stated earlier, SPs have not only to specify resource requirements but also constraints on the carbon emissions generated by the CPs while hosting their VDC. Specifically, green terms in the SLA specify the limit on carbon emissions that the CP is allowed to generate to accommodate the VDC request during a period of time called hereafter the *reporting period*. The reporting period can be for instance the a billing period [7].

To enforce green SLAs, the CP should compute the carbon footprint of each VDC request. To do so, we use two metrics: (1) carbon emission per unit of bandwidth (tonCO2/Mbps) and (2) carbon emission per core (tonCO2/Core). These metrics are chosen because the bandwidth and the CPU are the major factors that determine the power consumption in data centers and they are already considered in industry. For instance, Akamai reports annually its carbon emission in CO_2 per gigabyte of data delivered (tonCO2/Gbps), Verizon reports its carbon emissions per terabyte of transported data across its network.

As the carbon footprint is computed for each VDC, the SLA is enforced at the end of each reporting period. In case of violation of the green terms (i.e., the carbon footprint for the VDC is higher than the limit specified in the SLA), the CP is required to pay a penalty (a.k.a. credit). The penalty can a percentage of the SP's bill that can go up to 100% for some providers such as Rackspace [24]. It becomes then critical to design effective VDC embedding algorithms that minimize this penalty.

IV. Problem Formulation

In this section, we formally define the VDC embedding problem across multiple data centers as an Integer Linear Program (ILP). We assume that time is divided into slots. The metrics characterizing each data center (e.g., Power Usage Effectiveness (PUE), electricity price, availability of renewable power) are measured at the beginning of each time slot and are considered constant during the corresponding time slot. Moreover, we assume that the CP reports its carbon emissions periodically and every reporting period is T time slots. We denote by $T^k = [t_b^k, t_e^k]$ the k^{th} reporting period, where t_b^k and t_e^k are its beginning and end time slots, respectively.

The physical infrastructure is represented by a graph $G(V \cup W, E)$, where V denotes the set of data centers and W the set of nodes of the backbone network. The set of edges E represents the physical links in the backbone network. Each link is characterized by its bandwidth capacity $bw(e)$ and propagation delay $d(e)$.

A VDC request j is represented by a graph $G^j(V^j, E^j)$, its arrival time denote by t^j, and its lifetime T^j. Each vertex $v \in V^j$ corresponds to a VM, characterized by its CPU, memory and disk requirements. Each edge $e \in E^j$ is a virtual link that connects a pair of VMs, which is characterized by its bandwidth demand $bw(e)$ and propagation delay $d(e)$. We assume the revenue generated by VDC j, denoted by $\mathcal{R}^j$, to be proportional to the amount of CPU and bandwidth required by its VMs and links. The revenue generated by VDC j per time slot can be written as follows:

$$\mathcal{R}^j = \Big(\sum_{v \in V^j} (C^{cpu}(v) \times \sigma^{cpu}) + \sum_{e' \in E^j} bw(e') \times \sigma^b\Big) \quad (1)$$

where $C^{cpu}(v)$ is the CPU demand of VM v belonging to the VDC j, and σ^{cpu} and σ^b are unit price of CPU and bandwidth, respectively. Moreover, each VDC j may have a constraint on carbon emissions per reporting period T, which is defined by the variable c_j.

Furthermore, a VM $v \in V^j$ may have a location constraint. That is, it can only be embedded in a particular set of data centers. To model this constraint, we define a binary variable z_{ik}^j, indicating whether or not a VM k of VDC j can be embedded in a data center i.

The problem of embedding VDC requests in a distributed infrastructure of data centers should be solved dynamically over time. In fact, the decision of embedding VMs in different data centers is modified at the beginning of every time slot in such a way to follow the renewables. Thus, for each VDC request j, and during each time slot $t \in [t^j, t^j + T^j]$, the central controller should:

- Assign each VM $k \in V^j$ to a data center. Hence, we define the decision variable $x_{ik}^{j,t}$ as:

$$x_{ik}^{j,t} = \begin{cases} 1 & \text{If the VM } k \text{ of the VDC } j \text{ is assigned} \\ & \text{to data center } i \text{ during slot } t \\ 0 & \text{Otherwise.} \end{cases}$$

- Embed every virtual link either in the backbone network if it connects two VMs assigned to different data centers or within the same data center, otherwise. To do so, we define the virtual link allocation variable $f^t_{e,e'}$ as:

$$f^t_{e,e'} = \begin{cases} 1 & \text{If the link } e \in E \text{ is used to embed} \\ & \text{the virtual link } e' \in E^j \text{ during slot } t \\ 0 & \text{Otherwise.} \end{cases}$$

As an CP can reject a request due to shortage in resources or too tight constraints (delay, location). As such, we define a binary variable X_j, which indicates whether the VDC request j is accepted for embedding (i.e., $\sum_{t\in T^k}\sum_{i\in V}\sum_{k\in V^j} x^{j,t}_{ik} \geq 1$) or not.

Finally, the ultimate objective of the CP when embedding VDC requests during any reporting period T^k is to maximize its profit defined as the difference between the revenue (denoted by $\mathcal{R}_k$) and the total embedding cost plus penalty cost, which consists of the embedding cost in the data centers (denoted by $\mathcal{D}_k$), the migration cost (denoted by $\mathcal{M}_k$) the embedding cost in the backbone network $\mathcal{B}_k$ and the penalty cost $\mathcal{P}_k$. Hence, our problem can be formulated as an ILP with the following objective function:

$$Maximize \quad \mathcal{R}_k - (\mathcal{D}_k + \mathcal{B}_k + \mathcal{M}_k + \mathcal{P}_k) \tag{2}$$

Subject to:

$$x^{j,t}_{ik} \leq z^j_{ik},\ \forall k \in V^j, \forall i \in V, \forall t, \tag{3}$$

$$\sum_{i\in V} x^{j,t}_{ik} = X_j,\ \forall k \in V^j, \forall j \in Q_t, \forall t \tag{4}$$

$$\sum_{e'\in E^j} f^t_{e,e'} \times bw(e') \leq bw(e),\ \forall e \in E, \forall t \tag{5}$$

$$\sum_{e\in E} f^t_{e,e'} \times d(e) \leq d(e'),\ \forall e' \in E^j, \forall t \tag{6}$$

$$f^t_{e_1,e'} - f^t_{e_2,e'} = x^t_{dst(e_1)dst(e')} - x^t_{src(e_2)src(e')},$$
$$\forall e_1, e_2 \in E, dst(e_1) = src(e_2),\ \forall e' \in V^j, \forall t \tag{7}$$

where Q_t is the set of VDC requests being embedded during time slot t, $src(e)$ and $dst(e)$ denote the source and destination of link e, respectively. Equation (3) guarantees location constraint satisfaction. Equation (4) depicts that a VM is assigned to at most one data center. Equation (5) guarantees that link capacities are not exceeded in the backbone network, whereas (6) guarantees that delay requirements of virtual links are satisfied. Equation (7) denotes the flow continuity constraint.

The revenue for a reporting period T^k is given by:

$$\mathcal{R}_k = \sum_{t\in T^k}\sum_{j\in Q_t} \mathcal{R}^j \times X_j \tag{8}$$

Let us now focus on the expression of the embedding costs $\mathcal{D}_k$, $\mathcal{B}_k$, $\mathcal{M}_k$ and $\mathcal{P}_k$ in the data centers, the backbone network and penalty, respectively. Recall that these costs are part of the objective function.

- The cost of embedding in the data centers

In this work, we evaluate the request embedding cost in the data centers in terms of energy costs.

The total amount of consumed power in data center i is given by:

$$P^t_i = (P^t_{i,Net} + P^t_{i,Serv}) \times PUE^t_i \tag{9}$$

where $P^t_{i,Serv}$ and $P^t_{i,Net}$ are the power consumed by servers and network elements, respectively, and PUE^t_i is the power usage effectiveness of data center i during time slot t, which is used to compute the power consumed by supporting systems such as the cooling system. Note that this power consumption depends mainly on the local allocation scheme in each data center.

The mix of power used in data center i is given by:

$$P^t_i = P^t_{i,L} + P^t_{i,D} \tag{10}$$

where $P^t_{i,L}$ and $P^t_{i,D}$ denote, respectively, the consumed onsite renewable power and the amount of purchased power from the grid during time slot t. Note that $P^t_{i,L}$ should not exceed the amount of produced power, which is captured by $P^{j,t}_{i,L} \leq RN^t_i$, where RN^t_i is the amount of onsite renewable power generated in data center i, during time slot t, expressed in kW.

Hence, the total embedding cost in all data centers (expressed in \$) can be written as:

$$\mathcal{D}_k = \sum_{t\in T^k}\sum_{i\in V} P^t_{i,L} \times \eta_i + P^t_{i,D} \times \zeta^t_i \tag{11}$$

where η_i is the onsite renewable power cost in data center i (\$/kWh), ζ^t_i is the electricity price in data center i (\$/kWh).

- The cost of embedding in the backbone network

Virtual links between the VMs that have been assigned to different data centers should be embedded in the backbone network. We assume that it is proportional to their bandwidth requirements and the length of physical paths to which they are mapped. It is given by:

$$\mathcal{B}_k = \sum_{t\in T^k}\sum_{e'\in E^j}\sum_{e\in E} f^t_{e,e'} \times bw(e') \times \sigma_p \tag{12}$$

where σ_p is the cost incurred by the CP per unit of bandwidth allocated in the backbone network.

- The migration cost

Let's denote by t^- the previous time slot of time slot t. The migration cost is given by:

$$\mathcal{M}_k = \sum_{t\in T^k}\sum_{j\in(Q_{t^-}\cap Q_t)}\sum_{k\in V_j}\sum_{i\in V} mig^t_{k,i,j} \times m_{k,j} \tag{13}$$

where $m_{k,j}$ is the cost of migrating VM k of VDC j, which corresponds to the disruption in service that might occur when migrating the VM, and $mig^t_{k,i,j}$ is a binary variable that determines whether VM k of VDC j have been migrated to data center i from another data center at the beginning of time slot t, and defined as:

$$mig^t_{k,i,j} = \begin{cases} 1 & \text{If } x^{j,t}_{ik} = 1 \text{ and } x^{j,t-1}_{ik} = 0 \\ 0 & \text{Otherwise.} \end{cases}$$

Note that we assume that there is no cost for link migration as no transfer is needed.

- The penalty cost

The penalty is paid by the CP to the SP whenever the specified green SLA is not met. At the end of every reporting period T^k, the CP reports the carbon emission related to each VDC request j that has been embedded for the whole time period T^k or during a part of it. Since the carbon emissions are due to the power consumption, we can derive the carbon emission of every data center i during a time slot t, denoted by $\mathcal{C}_i^t$, as follows:

$$\mathcal{C}_i^t = P_{i,D}^t \times C_i \tag{14}$$

where $P_{i,D}^t$ denotes the amount of purchased power from the grid by data center i during time slot t and C_i is the carbon footprint per unit of power used from the grid in data center i expressed in tons of carbon per kWh (tonsCO2/kWh).

We derive the carbon emissions, in the entire infrastructure, due to the servers (denoted by $\mathcal{C}_{i,Serv}^t$) and the network (denoted by $\mathcal{C}_{Net}^t$), as follows:

$$\mathcal{C}_{Serv}^t = \frac{1}{|V|} \sum_{i \in V} \frac{\mathcal{C}_i^t \times P_{i,Serv}^t}{P_{i,Net}^t + P_{i,Serv}^t} \tag{15}$$

$$\mathcal{C}_{Net}^t = \frac{1}{|V|+1} \times (\sum_{i \in V} \frac{\mathcal{C}_i^t \times P_{i,Serv}^t}{P_{i,Net}^t + P_{i,Serv}^t} + \mathcal{C}_{Bckb}^t) \tag{16}$$

where $\mathcal{C}_{Bckb}^t$ is the carbon emission due to embedding virtual links in the backbone network. Without loss of generality, we assume that the carbon footprint per unit of bandwidth in the backbone network, denoted by C_b, is known. As such, $\mathcal{C}_{Bckb}^t$ is given by:

$$\mathcal{C}_{Bckb}^t = \sum_{e' \in E^j} \sum_{e \in E} f_{e,e'}^t \times bw(e') \times C_b \tag{17}$$

In this case, the average carbon emission rate of the CP per unit of VM during a reporting period T^k is given by:

$$\mathcal{C}_{CPU}^k = \frac{1}{t_e^k - t_b^k} \times \sum_{t \in [t_b^k, t_e^k]} \frac{\mathcal{C}_{Serv}^t}{\sum_{j \in Q_t} \sum_{v \in V^j} C^{cpu}(v)} \tag{18}$$

where Q_t is the set of VDC requests being embedded during time slot t and $C^{cpu}(v)$ is the capacity of VM v in terms of CPU units.

Similarly, the carbon emission rate per unit of bandwidth during a period T^k can be given as:

$$\mathcal{C}_{BW}^k = \frac{1}{t_e^k - t_b^k} \times \sum_{t \in [t_b^k, t_e^k]} \frac{\mathcal{C}_{Net}^t}{\sum_{j \in Q_t} \sum_{e \in E^j} bw(e)} \tag{19}$$

As such, the carbon emission related to a VDC request j during the period T^k, denoted by $\mathcal{C}_k^j$, can be given by:

$$\mathcal{C}_k^j = T_k^j \times ((\sum_{v \in V^j} C^{cpu}(v) \times \mathcal{C}_{CPU}^k) + (\sum_{e \in E^j} bw(e) \times \mathcal{C}_{BW}^k))$$

where T_k^j is the number of time slots of the period T^k during which VDC j was embedded.

Finally, a penalty is paid by the CP for an SP j at the end of the period T^k if the carbon emission for VDC j is above the limit specified in the SLA, i.e., $\mathcal{C}_k^j > c_j$, where c_j is the amount of carbon emission allowed by the SP for one reporting period T.

In the case where the CP has to pay a penalty, which is a fraction of the total bill during that period of time. Finally, the total penalty cost for a period T^k is given by:

$$\mathcal{P}_k = \sum_{j \in (\cup_{t \in T^k} Q_t)} (\mathcal{R}^j \times T_k^j) \times p, \text{ if } \mathcal{C}_k^j > c_j \tag{20}$$

where $p \in [0, 1]$ is the proportion of the SP's bill to be refunded by the CP in case of SLA violation.

The problem described above can be seen as a combination of the bin-packing problem and the multi-commodity flow problem, which are known to be $\mathcal{NP}$-hard. Therefore, we propose a simple yet efficient and scalable solution.

V. Green SLA opTimzER (GreenSlater)

Since the problem presented in the previous section is $\mathcal{NP}$-hard, we propose a greedy three-step approach. At the arrival a VDC request, the Central Controller first splits it into partitions such that the intra-partition bandwidth is maximized and the inter-partition bandwidth is minimized. It then uses an admission control algorithm that rejects VDCs with negative profit (i.e., the VDC cost is higher than the generated revenue). If the VDC is accepted, its partitions are embedded in different data centers. As the availability of renewables and electricity prices are variable over time, and the requests dynamically arrive and leave the system, we propose a reconfiguration algorithm, which migrates partitions from the data centers with no available renewables to those with available renewables. In the following, we present in details the proposed algorithms. Note that the partitioning aims at minimizing the backbone networks cost, while the reconfiguration minimizes the energy cost and limits the SLA violation by following the renewables, while taking into account the migration costs before migrating.

A. VDC Partitioning

Once received, the Central Controller divides the VDC request into partitions where the intra-partition bandwidth is maximized and the inter-partition bandwidth is minimized. Hence, each entire partition is then embedded in the same data center, which minimizes the inter-data center bandwidth. As the partitioning problem is $\mathcal{NP}$-hard [25], we use the Location Aware Louvain Algorithm (LALA), the partitioning algorithm used in [6]. LALA is a modified version of the Louvain Algorithm [26] that considers location constraints. The objective of the Louvain algorithm is to maximize the modularity, defined as an index between -1 and 1 that measures intra-partition density (i.e., the sum of the links' weights inside partitions) compared to inter-partition density (i.e., sum of the weights of links between partitions). In fact, graphs with high modularity have dense connections (i.e., high sum of weights) between the nodes within partitions, but sparse connections across partitions. Similar to the Louvain algorithm, the complexity of LALA is *O(n log n)* [26].

Algorithm 1 Admission Control Algorithm

IN: $predictionWdW$ // *the prediction window*
IN: $reconfigInterval$ // *the reconfiguration interval*
IN: vdc // *the VDC to embed*
$wdw \leftarrow min(predictionWdw, reconfigInterval)$
$possible \leftarrow possibleToEmbed(vdc)$
if $possible$ **then**
  $carbonRate \leftarrow getEstimationCarbonRate(wdw)$
  $carbonLimitRate \leftarrow vdc.carbonLimit/wdw$
  if carbonRate $\leq$ carbonLimitRate **then**
    Accept vdc
  else
    //Verify if profit can be made
    $estimatedCost \leftarrow estimatePowerCost(vdc)$
    if $revenue(vdc) \times (1 - refundFactor) - estimatedCost > 0$ **then**
      Accept vdc
    else
      Reject vdc
    end if
  end if
else
  Reject vdc
end if

Algorithm 2 Greedy VDC Partitions Embedding Across Data Centers

IN: $G(V \cup W, E)$, $G_M^j(V_M^j, E_M^j)$
for all $i \in V$ **do**
  $ToDC[i] \leftarrow \{\}$
end for
for all $v \in V_M^j$ **do**
  $S_v \leftarrow \{i \in V \ / \ i$ satisfies the location constraint$\}$
end for
for all $v \in V_M^j$ **do**
  $i \leftarrow s \in S_v$ with the smallest cost $getCost(s, v)$, and $LinksEmbedPossible(s, v) = true$
  if no data center is found **then**
    return $FAIL$
  end if
  $ToDC[i] \leftarrow ToDC[i] \cup \{v\}$
  for all $k \in N(v)$ **do**
    if $k \in ToDC[i]$ **then**
      $ToDC[i] \leftarrow ToDC[i] \cup \{e_{vk}\}$
    else
      if $\exists l \neq i \in V \ / \ k \in ToDC[l]$ **then**
        Embed e_{vk} in G using the shortest path
      end if
    end if
  end for
end for
return $ToDC$

B. Admission Control

When a VDC request is received, the Central Controller checks if the request will generate profit, in which case it is accepted, otherwise it is rejected. In some cases, a request with tight carbon constraints might result in high SLA violation penalties, which reduces the CP's profit. To address this issue, we propose an admission control algorithm (Algorithm 1). The idea is to estimate the available renewable power in the next prediction window and estimate carbon emission of the requested VDC. In this paper, we consider solar panels to generate the renewable power and we use a prediction model presented in [13].

First, the central controller checks whether it is possible to embed the VDC given the available resources and constraints of the VMs in the VDC. If the request is embeddable, the central controller computes an estimation for carbon emission for the request given the current power consumption and the predicted availability of renewables for the next prediction window. To do so, we propose to use a simple estimation algorithm, which computes the estimation of carbon emission per unit of VM and per unit of bandwidth in the next prediction window, and by the same derives the estimation of carbon emission of the given VDC request. The estimated carbon of the VDC request is then compared to the limit provided in the SLA of the VDC request. In case of SLA violation, the Central Controller checks whether profit can still be made even if there is a penalty to pay. If the profit is positive, the VDC request is accepted, otherwise it is rejected. It is worth noting that as the prediction window is limited compared to the lifetime of some of the VDCs (up to weeks for long lived VDCs), the decision of accepting might be biased as the short term forecasts can show high availability of renewables.

C. Partitions Embedding

Once a request $G^j(V^j, E^j)$ is partitioned, the resulting partitions that are connected through virtual links can be seen as a multigraph $G_M^j(V_M^j, E_M^j)$ where V_M^j is the set of nodes (partitions) and E_M^j is the set of virtual links connecting them. This multigraph is then embedded into the infrastructure, partition by partition, using Algorithm 2. For each partition $v \in V_M^j$, we first build the list of data centers that satisfy the location constraints of its VMs. The Central Controller queries the Local Controller of each data center s from the list to get the embedding cost of v. The cost is returned by the remote call $getCost(s, v)$.

The data center offering the lowest cost (provided by the procedure $getCost(s, v)$) and able to embed virtual links between v and all previously embedded partitions (denoted by $N(v)$) (verified by the function $LinksEmbedPossible(s, v)$) is then selected to host the partition. These virtual links are embedded in the backbone network using the shortest path algorithm.

This procedure is repeated until all partitions and virtual links that connect them are embedded into the distributed infrastructure. The complexity of embedding the whole multigraph is $O(|V_M^j| \times |V|)$, where $|V_M^j|$ is the number of partitions and $|V|$ is the number of data centers.

D. Dynamic Partition Relocation

As the the electricity price and the availability of renewables are variable over time, we propose a dynamic reconfiguration algorithm that optimizes VDC embedding over-time. The aim of the algorithm is to migrate partitions that have already been embedded in data centers which may run out of renewables towards data centers with available renewable power. The second criterion to perform a migration is to move partitions to locations where the electricity

Algorithm 3 Greedy Partition Migration Across Data Centers

1: **IN**: $predictionWdW$ // *the prediction window*
2: **IN**: $reconfigInterval$ // *the reconfiguration interval*
3: $wdw \leftarrow min(predictionWdW, reconfigInterval)$
4: **for all** $i \in V$ **do**
5: $Diff[i] \leftarrow EstimateRenewables(wdw, i) - FutureConsumption(wdw, i)$
6: **if** $Diff[i] < 0$ **then**
7: $part[i] \leftarrow$ list of partitions in i sorted by migration cost
8: **end if**
9: **end for**
10: **for all** $i \in V, Diff[i] < 0$ **do**
11: **while** $\exists\ k \in V, Diff[k] > 0$ **do**
12: $p \leftarrow part[i].first$
13: $D \leftarrow \{k \in V, Diff[k] > 0\}$
14: $done \leftarrow false$
15: **while** $!done$ && $D \neq \phi$ **do**
16: *//Take the data center with the minimum cost in the backbone network after migration*
17: $dest \leftarrow minBackboneCost(D)$
18: $Migrate(p, dest)$
19: **if** successful migration **then**
20: $done \leftarrow true$
21: Update $Diff[dest]$ and $Diff[i]$
22: **else**
23: $D \leftarrow D \backslash \{dest\}$
24: **end if**
25: **end while**
26: **end while**
27: **end for**

price is lower. We, hence, propose a migration algorithm (Algorithm 3) executed every τ hours (i.e., reconfiguration interval) by the central controller.

Data centers are first classified into two categories: sources and destinations. A data center is considered as a source if it has not enough renewable power to support its workload and hence we will have to resort to power from the grid. In this case, in a source data center, the difference between the estimated available renewable power and the estimated power consumption is negative (Line 5). Conversely, if a data has renewable power that exceeds its estimated power consumption, it is considered a destination since there is no need to reduce its workload and migrate VMs. In this case, a destination data center might be able to host more partitions if it has enough renewable power.

The idea is that partitions from source data centers should be migrated to destination data centers. To do so, the list of partitions in each source data center are sorted in increasing order of their migration cost (Line 7). For each partition, one destination data center that have a positive difference is chosen. The destination is chosen in a way that minimizes the inter-data center virtual link embedding cost after migration.

VI. Performance Evaluation

To evaluate the performance of Greenslater, we conducted several simulations using a realistic topology and real traces for electricity prices and renewable power availability. In the following, we first describe the simulation setting and we then present the results.

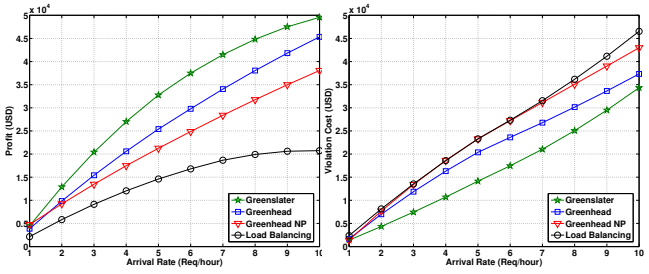

(a) Cumulative profit (b) SLA violation cost

Fig. 2: Impact of variable arrival rate λ ($P_{loc} = 0.05$, $T = 24$ hours, $\tau = 4$ hours)

A. Simulation Settings

For our simulations, we consider a physical infrastructure of 4 data centers located at four different states: New York, Illinois, California and Texas. The data centers are connected through the NSFNet topology as a backbone network, which includes 14 nodes. Each data center is connected to the backbone network through the closest node to its location. We assume all NSFNet links have a capacity of $100\,Gbps$. The traces of electricity prices and availability of renewable energy are provided by the US Energy Information Administration (EIA) [27]. The weather forecast is taken from the National Renewable Energy Laboratory [28] and the amount of power generated per square meter of solar panel from [29]. The carbon footprint per unit of power is provided by [30].
Similar to previous works [6], [14], VDCs are generated randomly according to a Poisson process with arrival rate λ and a lifetime following an exponential distribution with mean $1/\mu$. The number of VMs per VDC is uniformly distributed between 10 and 50. A pair of VMs belonging to the same VDC are directly connected with a probability 0.5 with a bandwidth demand uniformly distributed between 10 and $50\,Mbps$ and a delay uniformly distributed between 10 and 100 milliseconds. Each VM has a number of cores uniformly distributed between 1 and 4. Moreover, in each VDC, a fraction of VMs, denoted by $P_{loc} \in [0, 1]$, is assumed to have location constraints and thus cannot be migrated, i.e., it can only be embedded in a specific set of data centers. Each VDC comes with a carbon limit constraint specified in the Green SLA. This limit is assumed to be uniformly distributed between 5 and 20 $kgCO_2$ per day. Finally, we fixed the reporting period T to 24 hours. When the Green SLA is not satisfied, the CP refunds 50% of the SP's bill for that specific period of time. To assess the effectiveness of our proposal, we compare Greenslater to three solutions : (i) Greenhead [6], (ii) Greenhead with No Partitioning (NP) (i.e., each VM is considered as a single partition), and (iii) the load balancing approach for VDC embedding [22]. We consider five metrics: (i) the profit of the CP, which is the difference between revenue and the sum of operational costs (i.e., power cost, backbone network cost) and the Green SLA violation penalty, (ii) the acceptance ratio (defined as the ratio of embedded requests out of the total receive requests by the CP), (iii) the carbon footprint generated by the whole infrastructure, (iv) the

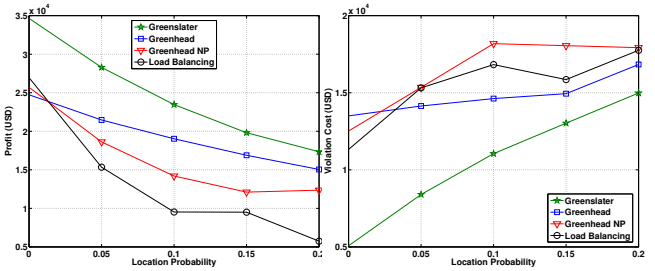

(a) Cumulative profit (b) SLA violation cost

Fig. 3: Impact of variable location probability P_{loc} ($\lambda = 4$ requests/hour, $T = 24$ hours, $\tau = 4$ hours)

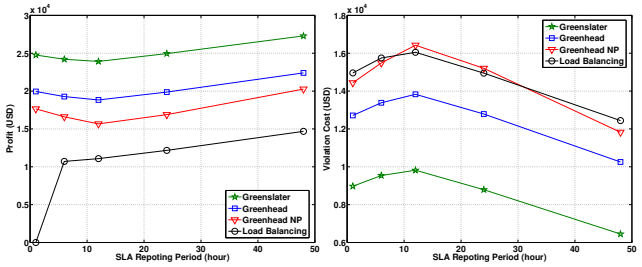

(a) Cumulative profit (b) SLA violation cost

Fig. 4: Impact of variable reporting period T ($\lambda = 4$ requests/hour, $P_{loc} = 0.05$, $\tau = 4$ hours)

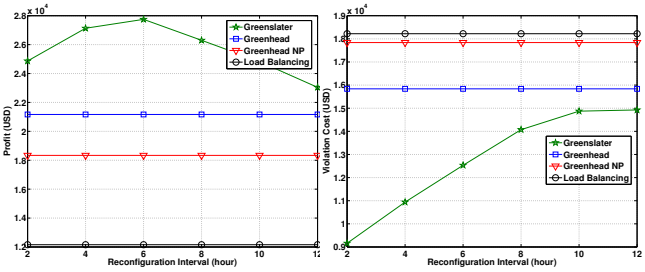

(a) Cumulative profit (b) SLA violation cost

Fig. 5: Impact of variable reconfiguration interval τ ($\lambda = 4$ requests/hour, $P_{loc} = 0.05$, $T = 24$ hours)

green power utilization and (v) the SLA violation penalty.

B. Simulation Results

In our simulations, we first study the impact of the different input parameters: the arrival rate λ, the fraction of location constrained VMs P_{loc} and the reporting period T on the system performance, using different values of the reconfiguration interval τ.

1) Impact of the Arrival Rate λ

Fig. 2 shows the impact of the arrival rate λ on both the achieved profit and SLA violation penalty, when $P_{loc} = 0.05$ (i.e., low constrained locations), $T = 24$ hours, and $\tau = 2$ hours. From this figure, we can notice that Greenslater outperforms other solutions, especially at high arrival rates (i.e., $\lambda \geq 3$). For small arrival rates (i.e., $\lambda \leq 2$), no considerable gain is observed as the number of requests being embedded is small. We can also observe that both the profit and SLA violation increase as the number of accepted requests increases. This is due to the fact that renewables are not enough to accommodate large numbers of VDCs, which leads to more power drawn from the electricity grid.

2) Impact of location probability constraint P_{loc}

Let us now study how location-constrained VMs may impact the results. To do so, we have varied P_{loc} between 0 and 0.2, and fixed the values of $\lambda = 4$ requests/hour, $T = 24$ hours and $\tau = 4$ hours. We can see in Fig. 3 that Greenslater outperforms the other solutions for all the values of P_{loc}. However, as P_{loc} increases, the profit drops for all approaches since more VMs must be located in specific data centers. This limits the possibility of migrating the partitions, which may run using power from the grid. It is clear that the gain achieved by Greenslater is higher when less location constraints are considered (i.e., low P_{loc}).

3) Impact of reporting period T

Fig. 4 shows the impact of reporting period T on both the achieved profit and the SLA violation cost. In this scenario, we vary T in $\{1, 6, 12, 24, 48\}$ hours, for fixed values of $\lambda = 4$ requests/hour, $P_{loc} = 0.05$ and $\tau = 4$ hours. Note that, in this case, the carbon constraint limit specified in the Green SLA is assumed to be uniformly distributed between 5 and 20 $kgCO_2$ per day, and is scaled down to the match the reporting period T. Again, Greenslater outperforms the baselines as it achieves higher profit and reduces the SLA violations costs. However, one can note that the profit is higher for long reporting periods (i.e., 24 and 48 hours) compared to short ones (i.e., 1,6 and 12 hours). The rational behind this is that for long reporting periods T, the CP has more time and more flexibility. In fact, the carbon footprint is computed as an average value over the whole period T. For small values of T, the CP does not have enough leverage since, in some data centers, VMs cannot be migrated even though renewables are available. This results in more frequent violation of the Green SLAs, which results in higher violations costs, as shown in Fig. 4(b), and thus lower profit (see Fig. 4(a)).

4) Impact of reconfiguration interval τ

We also study the impact of the reconfiguration interval τ on the profit and SLA violation penalty. We varied τ between 1 and 12 hours and fixed other variables ($\lambda = 4$ requests/hour, $P_{loc} = 0.05$ and $T = 24$ hours). The results are shown in Fig. 5. From this figure, we can see that the profit for Greenslater is a concave function of τ, where the maximum profit is obtained for $\tau_{opt} = 6$ hours in our case. In addition, the SLA violation penalty increases with τ, but remains low compared to the other solutions. In particular, for high values of τ, Greenslater gains decrease, since in this range of τ, the system configuration is not reoptimized to follow the renewables. Note that the variation of τ does not affect the performance of the other schemes, since they do not perform any migrations.

5) Summary of the results

To highlight the benefits of Greenslater over existing solutions, we plotted all the studied performance metrics (acceptance ratio, cumulative profit, utilization of renewable

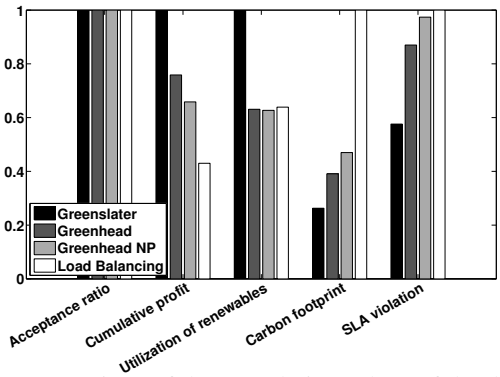

Fig. 6: Comparison of the cumulative values of the different metrics ($\lambda = 4$ requests/hour, $P_{loc} = 0.05$, $T = 24$ hours, $\tau = 4$ hours)

energy, carbon footprint and SLA violation cost) in Fig. 6. It is clear that Greenslater always achieves higher profit, ensures higher utilization of renewables and lower carbon footprint with minimum SLA violation. For instance, the gain in terms of profit provided by Greenslater is respectively around 33%, 53% and 120% compared to Greenhead, Greenhead NP and the Load Balancing approach.

VII. Conclusion

As the environmental impact of cloud infrastructures and services has become increasingly significant, governments and environmental organizations are urging SPs to require guarantees from their CPs that the carbon emission generated by the leased resources is limited. Hence, in this paper, we addressed the problem of including green constraints in the SLAs in order to cap the carbon emission of the resources allocated to each SP. We proposed Greenslater, a holistic framework that allows CPs to provision VDCs across a geographically distributed infrastructure with the goal of minimizing the operational costs and green SLA violation penalties. Results showed that Greenslater achieves high profit by minimizing operational costs and SLA violation penalties, while maximizing the utilization of the available renewable power.

Acknowledgments

This work was partially supported by the CNRS (France) WINDS project (Grant no. 25995), the NSERC (Canada) SAVI project (Grant no. NETGP394424-10), and the EU FP7 IRSES MobileCloud project (Grant no. 612212).

References

[1] ITU, "Toolkit on Environmental Sustainability for the ICT Sector (ESS)," 2012. [Online]. Available: http://www.itu.int/ITU-T/climatechange/ess/index.html

[2] KPMG, "Carbon footprint stomps on firm value," *KPMG International Report*, December 2012. [Online]. Available: http://goo.gl/migJnq

[3] "Carbon Disclosure Project website: www.cdp.com."

[4] IRRCi/Trucost, "Carbon risks and opportunities in the s&p 500," *Investor Responsibility Research Center Institute and Trucost Plc*, June 2009.

[5] Amazon Virtual Private Cloud, "http://aws.amazon.com/vpc/."

[6] A. Amokrane, M. F. Zhani, R. Langar, R. Boutaba, and G. Pujolle, "Greenhead: Virtual data center embedding across distributed infrastructures," *IEEE Transactions on Cloud Computing*, vol. 1, no. 1, pp. 36–49, 2013.

[7] Open Data Center Alliance, "Open Data Center Alliance Usage: Carbon Footprint Values," 2011. [Online]. Available: http://goo.gl/QcEfhG

[8] Report of the second meeting of the Cloud Selected Industry Group, Service Level Agreements expert subgroup, April 2013. [Online]. Available: http://bit.ly/KUhR8v

[9] G. Laszewski and L. Wang, "Greenit service level agreements," in *Grids and Service-Oriented Architectures for Service Level Agreements*, P. Wieder, R. Yahyapour, and W. Ziegler, Eds. Springer US, 2010, pp. 77–88.

[10] C. Bunse, S. Klingert, and T. Schulze, "GreenSLAs: Supporting Energy-Efficiency through Contracts," in *Energy Efficient Data Centers*, ser. Lecture Notes in Computer Science, 2012, vol. 7396.

[11] A. Galati, K. Djemame, M. Fletcher, M. Jessop, M. Weeks, S. Hickinbotham, and J. McAvoy, "Designing an SLA Protocol with Renegotiation to Maximize Revenues for the CMAC Platform," in *Web Information Systems Engineering Workshops*, 2013, pp. 105–117.

[12] C. Atkinson, T. Schulze, and S. Klingert, "Facilitating greener it through green specifications," *IEEE Software*, vol. 31, no. 3, May 2014.

[13] M. Haque, K. Le, I. Goiri, R. Bianchini, and T. Nguyen, "Providing green slas in high performance computing clouds," in *International Green Computing Conference (IGCC)*, 2013.

[14] M. F. Zhani, Q. Zhang, G. Simon, and R. Boutaba, "VDC Planner: Dynamic Migration-Aware Virtual Data Center Embedding for Clouds," in *IFIP/IEEE Integrated Network Management Symposium (IM 2013)*, May 2013.

[15] A. Qureshi, R. Weber, H. Balakrishnan, J. Guttag, and B. Maggs, "Cutting the electric bill for internet-scale systems," *SIGCOMM Comput. Commun. Rev.*, vol. 39, no. 4, pp. 123–134, Aug. 2009.

[16] Q. Zhang, Q. Zhu, M. F. Zhani, and R. Boutaba, "Dynamic Service Placement in Geographically Distributed Clouds," in *International Conference on Distributed Computing Systems (ICDCS)*, 2012.

[17] D. Hatzopoulos, I. Koutsopoulos, G. Koutitas, and W. Van Heddeghem, "Dynamic virtual machine allocation in cloud server facility systems with renewable energy sources," in *IEEE International Conference on Communications (ICC)*, 2013, pp. 4217–4221.

[18] P. X. Gao, A. R. Curtis, B. Wong, and S. Keshav, "It's not easy being green," in *ACM SIGCOMM*, 2012.

[19] Y. Guo, Y. Gong, Y. Fang, P. Khargonekar, and X. Geng, "Energy and network aware workload management for sustainable data centers with thermal storage," *IEEE Transactions on Parallel and Distributed Systems*, no. 99, pp. 1–1, 2013.

[20] J. He, X. Deng, D. Wu, Y. Wen, and D. Wu, "Socially-Responsible Load Scheduling Algorithms for Sustainable Data Centers over Smart Grid," in *IEEE International Conference on Smart Grid Communications (SmartGridComm)*, November 2012.

[21] Z. Abbasi, M. Pore, and S. Gupta, "Impact of workload and renewable prediction on the value of geographical workload management." in *Second International Workshop on Energy Efficient Data Centers (E2DC), held as a part of ACM eEnergy*, 2013.

[22] Y. Xin, I. Baldine, A. Mandal, C. Heermann, J. Chase, and A. Yumerefendi, "Embedding virtual topologies in networked clouds," in *International Conference on Future Internet Technologies*, ser. CFI, 2011, pp. 26–29.

[23] L. Wang, S. U. Khan, D. Chen, J. Kolodziej, R. Ranjan, C. zhong Xu, and A. Zomaya, "Energy-aware parallel task scheduling in a cluster," *Future Generation Computer Systems*, vol. 29, no. 7, 2013.

[24] S. A. Baset, "Cloud slas: Present and future," *SIGOPS Operating Syst. Rev.*, vol. 46, no. 2, pp. 57–66, July 2012.

[25] S. E. Schaeffer, "Graph clustering," *Computer Science Review*, vol. 1, no. 1, pp. 27 – 64, 2007.

[26] V. D. Blondel, J.-L. Guillaume, R. Lambiotte, and E. Lefebvre, "Fast unfolding of communities in large networks," *Journal of Statistical Mechanics: Theory and Experiment*, vol. 10, p. 8, Oct. 2008.

[27] U.S.Energy Information Administration. [Online]. Available: http://www.eia.gov

[28] National Renewable Energy Laboratory, last visited Feb. 2014. [Online]. Available: http://www.nrel.gov/gis/solar.html

[29] The Renewable Resource Data Center (RReDC) website, "http://www.nrel.gov/rredc/," 2012.

[30] Carbon Footprint Calculator, "http://www.carbonfootprint.com," 2012.