Green Cloud Framework For Improving Carbon Efficiency of Clouds

Saurabh Kumar Garg[1], Chee Shin Yeo[2] and Rajkumar Buyya[1]

[1] Cloud Computing and Distributed Systems Laboratory
Department of Computer Science and Software Engineering
The University of Melbourne, Australia
{sgarg, raj}@csse.unimelb.edu.au
[2] Distributed Computing Group
Computing Science Department
Institute of High Performance Computing, Singapore
yeocs@ihpc.a-star.edu.sg

Abstract. The energy efficiency of ICT has become a major issue with the growing demand of Cloud Computing. More and more companies are investing in building large datacenters to host Cloud services. These datacenters not only consume huge amount of energy but are also very complex in the infrastructure itself. Many studies have been proposed to make these datacenter energy efficient using technologies such as virtualization and consolidation. Still, these solutions are mostly cost driven and thus, do not directly address the critical impact on the environmental sustainability in terms of CO_2 emissions. Hence, in this work, we propose a user-oriented Cloud architectural framework, i.e. Carbon Aware Green Cloud Architecture, which addresses this environmental problem from the overall usage of Cloud Computing resources. We also present a case study on IaaS providers. Finally, we present future research directions to enable the wholesome carbon efficiency of Cloud Computing.

Keywords: Cloud Computing, Green IT, Resource Management

1 Introduction

Cloud Computing provides a highly scalable and cost-effective computing infrastructure for running IT applications such as High Performance Computing (HPC), Web and enterprise applications which require ever-increasing computational resources. The emergence of Cloud Computing has rapidly changed the paradigm of ownership-based computing approach to subscription-oriented computing by providing access to scalable infrastructure and services on-demand. The Cloud users can store, access, and share any amount of information online. Similarly, small and medium enterprises/organizations do not have to worry about purchasing, configuring, administering, and maintaining their own computing infrastructure. They can instead focus on improving their core competencies by exploiting a number of Cloud Computing benefits such as low cost, datacenter efficiencies, on-demand computing resources, faster and cheaper software development capabilities.

However, Clouds are essentially datacenters hosting application services offered on a subscription basis. They require high energy usage to maintain their operations. Today, a typical datacenter with 1000 racks needs 10 Megawatt of power to operate [19]. High energy usage results in high energy cost. Thus, for a datacenter, the energy cost is a significant component of its operating and upfront costs. In addition, in April 2007, Gartner estimated that the Information and Communication Technologies (ICT) industry generates about 2% of the total global CO_2 emissions, which is equal to the aviation industry [8]. According to a report published by the European Union [1], a decrease in emission volume of 15–30% is required before the year 2020 to keep the global temperature increase below 2^oC. Thus, the rapidly growing energy consumption and CO_2 emission of Cloud infrastructure has become a key environmental concern [20][4].

Hence, energy efficient solutions are required to ensure the environmental sustainability of this new computing paradigm. Up to now, as datacenters are the major elements of Cloud Computing resources, most solutions primarily focus on minimizing the energy consumption of datacenters which indirectly minimizes the CO_2 emission [2]. However, although such solutions can decrease the energy consumption to a great degree, they do not ensure the minimization of CO_2 emissions as a whole. For example, consider a Cloud datacenter which uses cheap energy generated by coal. The usage of such a datacenter will only increase CO_2 emissions.

Therefore, we propose a user-oriented Carbon Aware Green Cloud Architecture for reducing the carbon footprint of Cloud Computing in a wholesome manner without sacrificing the Quality of Service (QoS) (such as performance, responsiveness and availability) offered by multiple Cloud providers. Our architecture is designed such that it provides incentives to both users and providers to utilize and deliver the most "Green" services respectively. Our evaluation results in the context of IaaS Clouds show that a large amount of CO_2 savings can be gained using our proposed architecture. The contributions of this paper are:

- a novel Carbon Aware Green Cloud Architecture that aims to reduce CO_2 emissions without impacting the service performance; and
- a Carbon Efficient Green Policy (CEGP) for carbon-based scheduling that can reduce the carbon footprint of Cloud Computing by 25% compared to a basic Cloud resource management system.

2 Related Work

Most works improve the energy efficiency of Clouds by addressing the issue within a particular datacenter and not from the usage of Clouds as a whole. They focus on scheduling and resource management within a single datacenter to reduce the amount of active resources executing the workload [2]. The consolidation of Virtual Machines (VMs), VM migration, scheduling, demand projection, heat management, temperature aware allocation, and load balancing are used as basic techniques for minimizing energy consumption. Virtualization plays an important role in these techniques due to its several benefits such as consolidation, live migration and performance isolation.

Some works also propose frameworks to enable the energy efficiency of Clouds from user and provider perspectives. From the provider perspective, GreenCloud architecture [16] aims to reduce virtualized datacenter energy consumption by supporting optimized VM migration and VM placement. Similar work is presented by Lefevre et al. [14] who propose Green Open Cloud (GOC). GOC is designed for next generation Cloud datacenter that supports facilities like advance reservation. GOC aggregates the workload by negotiating with users so that idle servers can be switch-off longer.

Although these works maximize the energy efficiency of Cloud datacenters, they do not consider CO_2 emission which measures the environmental sustainibility of Cloud Computing. Even if a Cloud provider has used most energy efficient solutions for building his datacenter, it is still not assured that Cloud Computing will be carbon efficient. Greenpeace [10] indicates that current datacenters are really not environmentally friendly as Cloud providers are more concerned about reducing energy cost rather than CO_2 emission. For instance, Google Datacenter in Lenoir, NC, USA, uses 50.5% of dirty energy generated by coal. Thus, our previous work [7] proposes policies to simultaneously maximize the Cloud provider's profit and minimize the CO_2 emission of its non-virtualized datacenters. Le et al. [13] consider a similar multi-datacenter scenario, but with a different perspective of leveraging green energy by capping the brown energy. In contrast, here we propose an architectural framework which focuses on reducing the carbon footprint of Cloud Computing as a whole. Specifically, we consider all the elements of Cloud computing including Software, Platform, and Infrastructure as a Service. We also present a carbon aware policy for IaaS providers.

3 Carbon Aware Green Cloud Architecture

We propose Carbon Aware Green Cloud Architecture (Figure 1), which considers the goals of both users and providers while curbing the CO_2 emission of Clouds. Its elements include:

1. **Third Party:** Green Offer Directory and Carbon Emission Directory listing available green Cloud services and their energy efficiency respectively;
2. **User:** Green Broker accepting Cloud service requests (i.e. software, platform, or infrastructure) and selecting the most green Cloud provider; and
3. **Provider:** Green Middleware enabling the most carbon efficient operation of Clouds. The components of this middleware vary depending on the Cloud offerings (i.e. SaaS, PaaS, or IaaS).

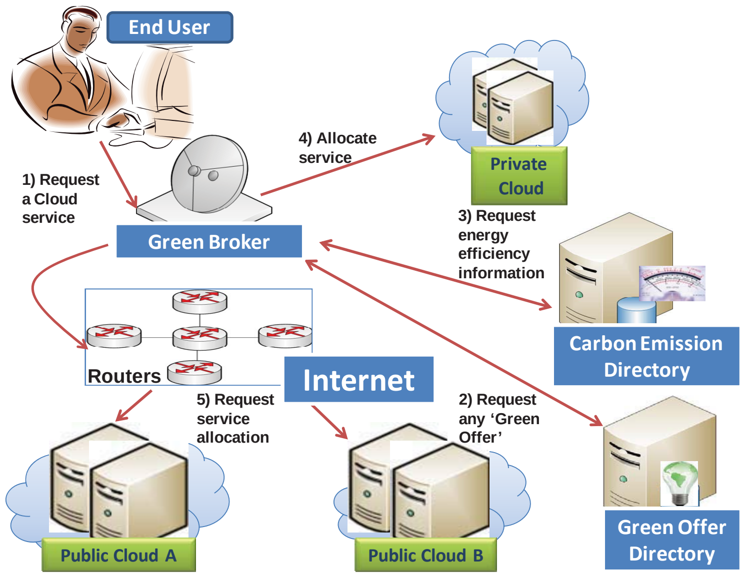

Fig. 1. Carbon Aware Green Cloud Architecture

3.1 Third Party: Green Offer Directory and Carbon Emission Directory

We propose two new elements, i.e. Green Offer Directory and Carbon Emission Directory, which are essential to enforce the green usage of Cloud Computing. Governments have already introduced energy ratings for datacenters and various laws to cap the energy usage of these datacenters [12][22]. There is also increasing awareness on the impact of greenhouse gases on climate change [10]. Therefore, users will likely prefer using Cloud services of providers which ensure the minimum carbon footprint. Cloud providers can also use these directories as an advertising tool to attract more users. For instance, Google has released the energy efficiency of its datacenters [17]. Hence, the introduction of such directories is practical in the current context of Cloud Computing.

Cloud providers register their services in the form of 'Green Offers' to a Green Offer Directory which is accessed by Green Broker. These offers consist of the type of service provided, pricing, and time when it can be accessed for the least CO_2 emission. The Carbon Emission Directory maintains data related to the energy efficiency of Cloud services, which include the Power Usage Effectiveness (PUE) and cooling efficiency of Cloud datacenters which are providing the service, network cost, and CO_2 emission rate of electricity. Hence, Green Broker can get the current status of energy parameters for using various Cloud services from Carbon Emission Directory.

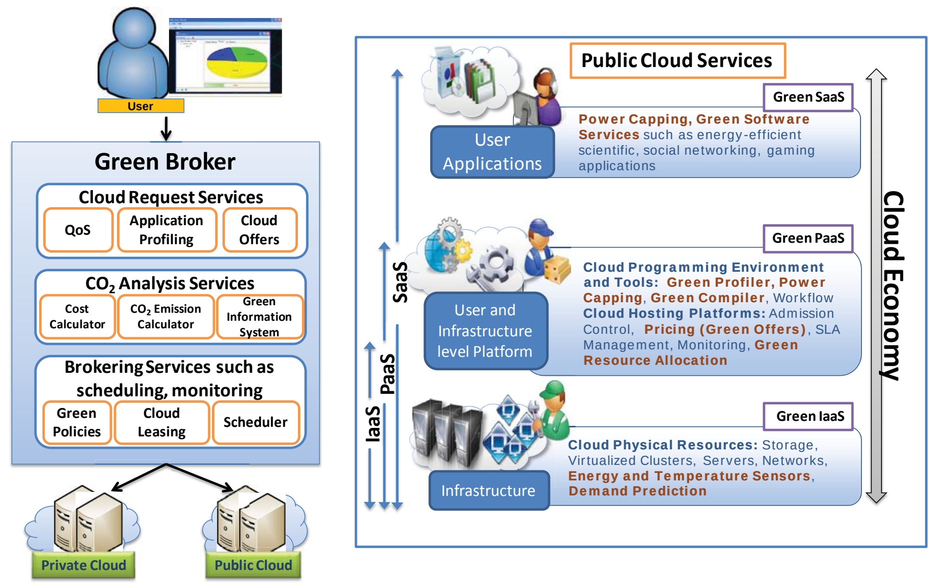

Fig. 2. (a) Green Broker and (b) Green Middleware components for each Cloud service (SaaS, PaaS, and IaaS)

3.2 User: Green Broker

Green Broker (Figure 2) has similar responsibility as a typical Cloud broker, i.e. to lease Cloud services on behalf of users and schedule their applications. Its first layer comprises Cloud request services that analyze the requests and their QoS requirements. Its second layer calculates the cost and carbon footprint of leasing particular Cloud services based on information about various Cloud offerings and current CO_2 emission factors obtained from Green Offer Directory and Carbon Emission Directory respectively. With these calculations, Green Policies make the decisions of leasing Cloud services. If no exact match is found for a request, alternate 'Green Offers' are suggested to users by Cloud Request Services.

The carbon footprint of a user request depends on the type of Cloud service it requires, i.e. SaaS, PaaS and IaaS, and is computed as the sum of CO_2 emission due to data transfer and service execution at datacenter. SaaS and PaaS requests use CO_2 emission per second (CO2PS) to reflect long term usage, while IaaS request uses CO_2 emission as data transfer is mostly once.

– **SaaS and PaaS Request (CO_2 emission per second):**

$$CO2PS_{SaaS/PaaS} = (r_{dT}^{CO2} E_{dT} \times a_{dT}) + (r^{CO2} \times \frac{1}{DCiE} \times E_{serv}) \quad (1)$$

where r_{dT}^{CO2} is the CO_2 emission rate per joule of energy spent from the user's machine to the datacenter, E_{dT} is the per-bit energy consumption of data transfer, a_dT is the data bits transferred per second, r^{CO2} is the CO_2 emission rate where the datacenter is located, $DCiE$ is the power efficiency of the datacenter defined as the fraction of total power dissipated that is used for IT resources, and E_{serv} is the energy spent per second by the server for executing the user's request. The total power dissipated by a Cloud provider is used not only for computers, but also for other purposes, including power conditioning, HVAC (Heating, Ventilating, and Air Conditioning), lighting, and wiring [9]. Therefore, DCiE is the most appropriate parameter for selecting Cloud providers.

– **IaaS Request ($\mathbf{CO_2}$ emission):**

$$CO2_{IaaS} = (r^{CO2}_{dT} E_{dT} \times IOdata) + (r^{CO2} \times \frac{1}{DCiE} \times E_{serv} \times Vtime) \quad (2)$$

where $IOdata$ is the data transferred to run application on VM leased from Clouds and $Vtime$ is the time for which VM is active.

3.3 Provider: Green Middleware

To support carbon aware Cloud Computing, a Cloud provider must implement "Green" conscious middleware at various layers depending on the type of Cloud service offered (SaaS, PaaS, or IaaS) (Figure 2) as follows:

– **SaaS Level:** SaaS providers mainly offer software installed in their own datacenters or resources leased from IaaS providers. Therefore, they require Power Capping component to limit the usage of software services by each user. This is especially important for social networking and game applications where users become completely unaware of their actions on environmental sustainability. SaaS providers can also offer Green Software Services deployed on carbon efficient datacenters with less replications.
– **PaaS Level:** PaaS providers in general offer platform services for application development and their deployment. Thus, to ensure energy efficient development of applications, relevant components such as Green Compiler to compile applications with the minimum carbon footprint and carbon measuring tools for users to monitor the carbon footprint of their applications. For example, JouleSort [19] is a Green Profiler providing energy efficiency benchmarks to measure the energy required to perform an external sort.
– **IaaS level:** IaaS providers play the most crucial role in the success of Green Cloud Architecture since IaaS not only offers independent infrastructure services, but also support other services (SaaS and PaaS) offered by Clouds. They use the latest technologies for IT and cooling systems to have the most energy efficient infrastructure. By using virtualization and consolidation, the energy consumption is further reduced by switching off unutilized servers. Energy and Temperature Sensors are installed to calculate the current energy efficiency of each IaaS provider and their datacenters. This information is advertised regularly by Cloud providers in the Carbon Emission Directory. Various green scheduling and resource provisioning policies will ensure

minimum energy usage. In addition, IaaS providers can design attractive 'Green Offers' and pricing schemes providing incentives for users to use their services during off-peak or maximum energy efficiency hours.

4 Case Study: IaaS Cloud

To illustrate the effectiveness of our proposed architecture in reducing the energy and CO_2 emissions across the entire Cloud infrastructure in a unified manner, we present a simple scenario focussed on IaaS. It considers multiple IaaS providers offering computational resources to run HPC jobs. A user request consists of an application, its estimated length in time, the deadline to complete execution, and the number of resources required. Requests are submitted to Green Broker which interprets and analyzes the service requirements before deciding where to execute them.

Cloud datacenters have different CO_2 emission rates and energy costs based on their locations. Each datacenter updates this data to Carbon Emission Directory for facilitating carbon efficient scheduling. For this study, we consider three CO_2 emission related parameters: CO_2 emission rate (kg/kWh) (r_i^{CO2}), average DCiE ($Ieff_i$), and VM power efficiency ($VMeff_i$). The VM power efficiency is the amount of power dissipated by fully active VM running at maximum utilization level [3]. In Green Offer Directory, IaaS providers specify the maximum number of VMs that can be initiated at a particular time for achieving the highest energy efficiency due to the variation in datacenter efficiency with time and load [18] and power capping technologies used within the datacenter [15].

5 Carbon Efficient Green Policy (CEGP)

We develop Carbon Efficient Green Policy (CEGP) for Green Broker to periodically select the Cloud provider with the minimum carbon footprint and initiate VMs to run the jobs (Algorithm 1). Based on user requests at each scheduling interval, Green Broker obtains information from Carbon Emission Directory about the current CO_2 emission related parameters of providers as described in Section 4 (Line 2). The QoS requirements of a job j is defined in a tuple (d_j, n_j, e_j, f_j^m), where d_j is the deadline to complete job j, n_j is the number of CPUs required for job execution, and e_j is the job execution time when operating at the CPU frequency f_j^m (Line 3).

CEGP then sorts the incoming jobs based on Earliest Deadline First (EDF) (Line 4), before sorting the Cloud datacenters based on their carbon footprint (Line 5). CEGP schedule jobs to IaaS Clouds in a greedy manner to reduce the overall CO_2 emission. For IaaS providers, CEGP uses three main factors to calculate the CO_2 emission: CO_2 emission rate, DCiE, and CPU power efficiency. The carbon footprint of an IaaS Cloud i is given by: $r_i^{CO2} \times \frac{1}{Ieff_i} \times \frac{1}{VMeff_i}$ where $VMeff_i$ can be calculated by Cloud providers based on the proportion of resources on a server utilized by the VM using tools such as PowerMeter [3]. If a VM consumes the power equivalent to a processor running at f_i frequency

```
1  while current_time < next_schedule_time do
2      RecvCloudPublish(P);
       //P contains information of Cloud datacenters
3      RecvJobQoS(Q);
       //Q contains information of Cloud users
4  Sort jobs in ascending order of deadline;
5  Sort datacenters in ascending order of r_i^{CO2} × 1/Ieff_i × 1/VMeff_i;
6  foreach job j ∈ RecvJobQoS do
7      foreach datacenter i ∈ RecvCloudPublish do
8          if isInitiatedVM(i) then
9              if MaxIniVMlimitReached(i) then
10                 Try to schedule the job j on already initiated VMs;
11                 if job j is missing deadline then
12                     continue;
13                 break;
14             else
15                 InitiateVM(i) and schedule job j;
16                 break;
```

Algorithm 1: Carbon Efficient Green Policy (CEGP)

level, then we can use the following power model [5][23] to calculate its power efficiency: $\beta_i + \alpha_i(f_i)^3$, where β_i is the static power dissipated by the CPU and α_i is the proportionality constant. Therefore, the approximate energy efficiency of VM is: $VMeff_i = |\frac{f_i}{\beta_i+\alpha_i(f_i)^3}|$. If job j executes at CPU frequency f, then its CO_2 emission will be the minimum when it is allocated to the datacenter with the minimum CO_2 emission rate r_i^{CO2}, maximum DCiE value $Ieff_i$, and maximum CPU power efficiency $VMeff_i$. CEGP then assigns jobs to VMs initiated on each Cloud datacenter according to this ordering (Line 6–16).

6 Performance Evaluation and Results

We use the Lawrence Livermore National Laboratory (LLNL) Thunder trace from Feitelson's Parallel Workload Archive (PWA) [6] with the highest resource utilization of 87.6% to ideally model a heavy HPC workload scenario. The trace contains the submit time, requested number of CPUs, and actual runtime of jobs. We use a methodology proposed by Irwin et al. [11] to synthetically assign deadlines through two classes, namely Low Urgency (LU) and High Urgency (HU). We set LU jobs to have a deadline mean of 12, which is 3 times longer than HU jobs with a deadline mean of 4. The arrival sequence of jobs from the HU and LU classes is randomly distributed.

Provider Configuration: We model 8 different IaaS providers with different configurations as listed in Table 1. Power parameters (i.e. CPU power factors and frequency level) of the CPUs at different datacenters are derived from Wang and Lu's work [23]. Green Broker uses CEGP to schedule jobs periodically at a scheduling interval of 50 seconds, which is to ensure that Green Broker can receive at least one job in every scheduling interval. The DCiE value of Cloud datacenters is randomly generated using a uniform distribution between [0.33, 0.80] as indicated in the study conducted by Greenberg et al. [9].

Table 1. Characteristics of Cloud datacenters

Location of Cloud Datacenter	CO_2 Emission Rate (kg/kWh) [a]	CPU Power Factors		CPU Frequency Level
		β	α	f_i
New York, USA	0.389	65	7.5	1.8
Pennsylvania, USA	0.574	75	5	1.8
California, USA	0.275	60	60	2.4
Ohio, USA	0.817	75	5.2	2.4
North Carolina, USA	0.563	90	4.5	3.0
Texas, USA	0.664	105	6.5	3.0
France	0.083	90	4.0	3.2
Australia	0.924	105	4.4	3.2

[a] CO_2 emission rates are drived from a US Department of Energy (DOE) document [21] (Appendix F-Electricity Emission Factors 2007).

Experimental Scenarios: We compare the carbon efficiency of CEGP with a performance-based scheduling algorithm (Earliest Start Time (EST)) using two metrics: average energy consumption and CO_2 emissions. EST schedules jobs to the datacenter where jobs can start as earliest as possible with the least waiting time. The average energy consumption shows the amount of energy saved by our green framework using CEGP compared to an existing approach using EST which just focus on performance, whereas the average CO_2 emission shows its corresponding environmental impact. We examine two experimental scenarios: 1) comparison of CEGP with EST and 2) effect of relationship between CO_2 emission rate and datacenter power efficiency $DCiE$. The first scenario demonstrates how our proposed architecture can achieve higher carbon efficiency. The second scenario reveals how the relationship between CO_2 emission rate and DCiE can affect the achievement of carbon efficiency. Hence, we consider two types of relationship between CO_2 emission rate and DCiE: 1) datacenter with the highest CO_2 emission rate has the highest DCiE (HH) and 2) datacenter with the highest CO_2 emission rate has the lowest DCiE (HL). We generate 8 DCiE values using uniform distribution between [0.33, 0.80] and assign them to the 8 datacenters to achieve HH and HL configurations accordingly.

6.1 Comparison of CEGP with Performance-based Algorithm (EST)

We compare CEGP with EST for datacenters with HH configuration. The effect of job urgency on energy consumption and CO_2 emission is prominent. As the percentage of HU jobs with more urgent (shorter) deadlines increases, the energy consumption (Figure 3(a)) and CO_2 emission (Figure 3(b)) also increase due to more urgent jobs running on datacenters with lower DCiE value and at the highest CPU frequency to avoid deadline violations.

It is clear that our proposed architecture using CEGP (EDF-CEGP) can reduce up to 23% of the energy consumption (Figure 3(a)) and 25% of the CO_2 emission (Figure 3(b)) compared to an existing approach using EST (EDF-EST) across all datacenters. CEGP is also able to complete very similar amount of workload[3] as EST (Figure 3(c)), but with much less energy consumption

[3] workload=$\sum$(job execution time $\times$ number of required processors)

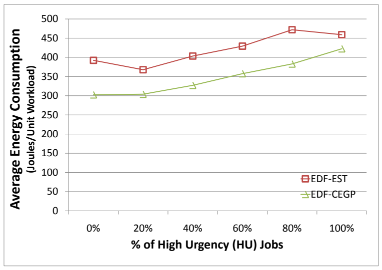

(a) Energy Consumption

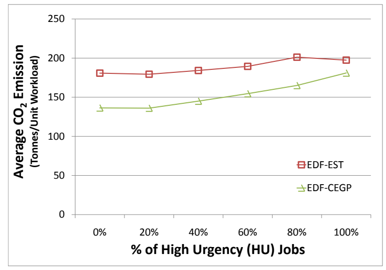

(b) CO_2 Emission

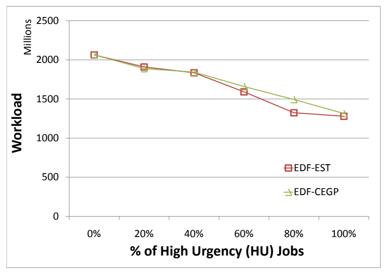

(c) Workload Executed

Fig. 3. Comparison of CEGP with performance-based algorithm (EST)

and CO_2 emission. This highlights the importance of considering the DCiE and CO_2 emission related factors in achieving the carbon efficient usage of Cloud Computing. In particular, CEGP can reduce energy consumption (Figure 3(a)) and CO_2 emission (Figure 3(b)) even more when there are more LU jobs with less urgent (longer) deadline.

6.2 Effect of Relationship between CO_2 Emission Rate and Datacenter Power Efficiency *DCiE*

This experiment analyzes the impact of different configurations (HH and HL) of datacenters with respect to CO_2 emission rate and datacenter power efficiency *DCiE* based on 40% of high urgency jobs.

In both HH and HL configurations, CEGP reduces CO_2 emission and energy consumption between 23% to 25% (Figure 4(a) and 4(b)). Therefore, we infer that for other configurations, we will also achieve similar carbon efficiency in Cloud Computing by using CEGP. Moreover, in Figure 4(a), there is a decrease in energy consumption of all the Cloud datacenters from HH to HL configuration by using EST, while there is almost no corresponding decrease by using CEGP. This shows that how important is the consideration of global factors such as DCiE and CO_2 emission rate in order to improve the carbon footprint of Cloud Computing.

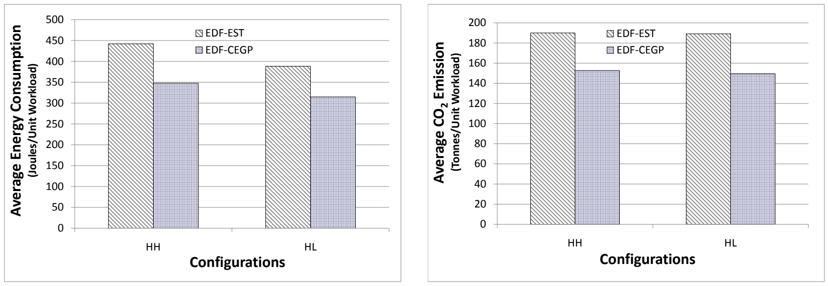

(a) Energy Consumption

(b) CO_2 Emission

Fig. 4. Effect of relationship between CO_2 emission rate and DCiE

7 Conclusion

In this paper, we present a Carbon Aware Green Cloud Architecture to improve the carbon footprint of Cloud Computing taking into account its global view. Our architecture is designed such that it provides incentives to both users and providers to utilize and deliver the most "Green" services respectively. Therefore, it embeds components such as Green broker from user side to ensure the execution of their applications with the minimum carbon footprint. Similarly, from provider side, we propose features for next generation Cloud providers who will publish the carbon footprint of their services in public directories and provide 'Green Offers' to minimize their overall energy consumption. We also propose a Carbon Efficient Green Policy (CEGP) for Green broker which schedules user application workload with urgent deadline on Cloud datacenters with more energy efficiency and low carbon footprint.

Further, the simulation-based evaluation of our architecture is done in multiple IaaS Cloud provider scenario. We compare two scheduling approaches to prove how our proposed architecture helps in improving carbon and energy footprint of Cloud Computing. Performance evaluation results show how our proposed architecture using a Green Policy CEGP can save up to 23% energy while improving the carbon footprint by about 25%. Therefore, these promising results show that by using our architectural framework carbon footprint and energy consumption of Cloud Computing can be improved.

In the future, we will investigate different 'Green Policies' for Green broker and also how Cloud providers can design various 'Green Offers' based on their internal power efficiency techniques such as VM consolidation and migration. We will also conduct experiments for our architecture using real Clouds.

References

1. Baer, P.: Exploring the 2020 global emissions mitigation gap. `http://www.ippr.org/uploadedFiles/globalclimatenetwork/Exploring_the_Mitigation_Gap[1].pdf` (Dec 2008)

2. Beloglazov, A., Buyya, R., Lee, Y., Zomaya, A.: A Taxonomy and Survey of Energy-Efficient Data Centers and Cloud Computing Systems. Advances in Computers, M. Zelkowitz (editor). Elsevier, San Francisco, USA (2011)
3. Bohra, A., Chaudhary, V.: Vmeter: Power modelling for virtualized clouds. In: Proc. of 24th IEEE IPDPS Workshops. Atlanta, USA (2010)
4. Cameron, K.: Trading in Green IT. Computer 43(3), 83–85 (2010)
5. Chen, Y., et al.: Managing server energy and operational costs in hosting centers. ACM SIGMETRICS Performance Evaluation Review 33(1), 303–314 (2005)
6. Feitelson, D.: Parallel workloads archive. http://www.cs.huji.ac.il/labs/parallel/workload (2011)
7. Garg, S., Yeo, C., Anandasivam, A., Buyya, R.: Environment-conscious scheduling of HPC applications on distributed cloud-oriented data centers. Journal of Parallel and Distributed Computing 71(6), 732–749 (2011)
8. Gartner: Gartner Estimates ICT Industry Accounts for 2 Percent of Global CO2 Emissions. http://www.gartner.com/it/page.jsp?id=503867 (Apr 2007)
9. Greenberg, S., et al.: Best practices for data centers: Results from benchmarking 22 data centers. In: ACEEE Summer Study on Energy Efficiency in Buildings (2006)
10. Greenpeace International: Make IT green: Cloud computing and its contribution to climate change (2010)
11. Irwin, D., Grit, L., Chase, J.: Balancing risk and reward in a market-based task service. In: Proc. of 13th IEEE HPDC. Honolulu, USA (2004)
12. Kurp, P.: Green computing. Commun. ACM 51, 11–13 (2008)
13. Le, K., et al.: Managing the cost, energy consumption, and carbon footprint of internet services. ACM SIGMETRICS Perf. Eval. Review 38(1), 357–358 (2010)
14. Lefèvre, L., Orgerie, A.: Designing and evaluating an energy efficient Cloud. The Journal of Supercomputing 51(3), 352–373 (2010)
15. Lefurgy, C., Wang, X., Ware, M.: Power capping: a prelude to power shifting. Cluster Computing 11(2), 183–195 (2008)
16. Liu, L., et al.: GreenCloud: a new architecture for green data center. In: Proc. of 6th International Conference on Autonomic Computing. Barcelona, Spain (2009)
17. Miller, R.: Google: Raise Your Data Center Temperature. http://www.datacenterknowledge.com/archives/2008/10/14/google-raise-your-data-center-temperature (Oct 2008)
18. Patel, C., et al.: Energy Aware Grid: Global Workload Placement based on Energy Efficiency. Technical Report HPL-2002-329, HP Labs, Palo Alto (2002)
19. Rivoire, S., Shah, M.A., Ranganathan, P., Kozyrakis, C.: Joulesort: a balanced energy-efficiency benchmark. In: Proc. of ACM SIGMOD. Beijing, China (2007)
20. Tomlinson, B., Silberman, M.S., White, J.: Can More Efficient IT Be Worse for the Environment? Computer 44, 87–89 (2011)
21. U.S. DOE: Voluntary Reporting of Greenhouse Gases: Appendix F. Electricity Emission Factors. http://www.eia.doe.gov/oiaf/1605/pdf/Appendix%20F_r071023.pdf (2007)
22. U.S. EPA: Report to Congress on Server and Data Center Energy Efficiency, Public Law 109-431 (Aug 2007)
23. Wang, L., Lu, Y.: Efficient Power Management of Heterogeneous Soft Real-Time Clusters. In: Proc. of 29th IEEE RTSS. Barcelona, Spain (2008)