

# Energy Efficient Service Delivery in Clouds in Compliance with the Kyoto Protocol

Drazen Lucanin<sup>1,2</sup>, Michael Maurer<sup>1</sup>, Toni Mastelic<sup>1</sup>, and Ivona Brandic<sup>1</sup>

<sup>1</sup> Vienna Univ. of Technology, Vienna, Austria  
{drazen,maurer,toni,ivona}@infosys.tuwien.ac.at  
<sup>2</sup> Ruder Boskovic Inst., Zagreb, Croatia

**Abstract.** Cloud computing is currently revolutionizing the ICT landscape by providing novel, scalable and cost & energy efficient models for on-demand access to geographically distributed computing resources. As well as any other industry branch (e.g. manufacturing, civil engineering) the ICT industry - Cloud infrastructures such as data centers in particular – are responsible for considerable amounts of  $CO_2$ . Thus, the operators of data centers will very soon be faced with legislative guidelines defining restrictions on  $CO_2$  emissions. One of the most prominent guidelines is the Kyoto protocol, defining particular  $CO_2$  caps at different organizational levels (level of a country, industry branch, etc.) Although, there is a large body of work done in the area of modeling and implementing energy efficient data centers, there is yet very little work done in defining a flexible  $CO_2$ -trading model that will represent the cornerstone for the implementation of next-generation data centers. In this paper we present a first attempt in modeling data centers in compliance with the Kyoto protocol. Thereby, we discuss a novel approach for trading of  $CO_2$  caps across data centers according the limitations that each data center has to comply with. Thus,  $CO_2$  caps can be embedded into Service Level Agreements and traded like other computing commodities (e.g. computational power, storage). Our model represents the foundation for the implementation of next-generation schedulers and pricing models that support Kyoto-compliant  $CO_2$  trading schemes.

## 1 Introduction

With the global advent of cloud, grid, cluster computing and increasing needs for large data centers to run these services, the environmental impact of large-scale computing paradigms is becoming a global problem. The energy produced to power the ICT industry (and data centers constitute its major part) is responsible for 2% of all the carbon dioxide equivalent ( $CO_2e$  – greenhouse gases normalized to carbon dioxide by their environmental impact) emissions [6], thus accelerating global warming [11].

Cloud computing facilitate users to buy computing resources from a cloud provider and specify the exact amount of each resource that they expect through a Service Level Agreement (SLA). The cloud provider then honors this agreement by providing the promised resources to avoid agreement violation penalties

(and to keep the customer satisfied to continue doing business). However, cloud provider are usually faced with the challenge of satisfying promised SLAs and at the same time not wasting their resources [14].

In order to fight global warming United Nations Framework Convention on Climate Change (UNFCCC or FCCC) agreed on the Kyoto protocol. The goal is to achieve global stabilisation of greenhouse gas concentrations in the atmosphere at a level that would prevent dangerous anthropogenic interference with the climate system [13]. Keeping track of and reducing  $CO_2e$  emissions is becoming more and more relevant after the ratification of the Kyoto protocol. Thereby, the protocol defines control mechanisms to reduce  $CO_2e$  emissions by basically setting a market price for such emissions. Currently, flexible models for  $CO_2e$  trading are developed at different organizational and political level as for example at the level of a country, industry branch, or a company.

On the one hand, there is large body of work done in facilitating energy efficient management of data centers as for example represented in [8] where current state of formal energy efficiency control in cloud computing relies on monitoring power usage efficiency (PUE) and the related family of metrics developed by the Green Grid Consortium. Another example is discussed in [15] where an economical incentive is presented to promote greener cloud computing policies. On the other hand, there are several mature models for trading of  $CO_2e$  caps in various industrial branches, as for example in the oil industry [17]. Surprisingly, to the best of our knowledge there is no existing related work about the application of Kyoto protocol to energy efficient modeling of data centers and Cloud infrastructures.

In this paper we propose an  $CO_2e$  trading model for transparent scheduling of resources in cloud computing following the guidelines of the Kyoto protocol [17]. First, we present a conceptual model for  $CO_2e$  trading compliant to the Kyoto protocol's emission trading scheme. We consider an *emission trading market (ETM)* where *credits for emission reduction (CERs)* are traded between data centers. Based on the positive or negative *CERs* applications are either in- or outsourcing in order to achieve a  $CO_2e$  equilibrium. Thereby, applications are in- and outsourced following the (i) energy costs, (ii)  $CO_2e$  costs and (iii) service costs. We propose a *wastage - penalty* model that can be used as a basis for the implementation of Kyoto protocol-compliant scheduling and pricing models. Finally, we discuss potential uses of the model as an optimisation heuristic in the resource scheduler.

The main contribution of the paper are (1) definition of the conceptual *emission trading market (ETM)* for the application of Kyoto protocol for the energy efficiency management in Clouds (2) definition of a *wastage - penalty* model for trading of *credits for emission reduction (CERs)* (3) discussion on how the presented *wastage - penalty* model can be used for the implementation of next generation Kyoto protocol compliant energy efficient schedulers and pricing models.

The paper is structured as follows: Section 2 discusses related work. Section 3 gives some background as to why cloud computing might become subject to the

Kyoto protocol. Section 4 presents our model in a general  $CO_2e$ -trading cloud scenario, we then go on to define a formal model of individual costs to find a theoretical balance and discuss the usefulness of such a model as a scheduling heuristic. Section 5 concludes the paper and identifies possible future research directions.

## 2 Related work

As our aim is to enable energy efficiency control in the cloud resource scheduling domain, there are two groups of work related to ours that deal with the problem:

1. scheduling algorithms - resource allocation techniques, from which energy cost optimisation is starting to evolve
2. energy efficiency legislation - existing rules, regulations and best behaviour suggestions that are slowly moving from optimising the whole data center efficiency towards optimising its constituting parts

We will examine each of these two groups separately now.

### 2.1 Scheduling algorithms

There already exist cloud computing energy efficient scheduling solutions, such as [21, 20] which try to minimize energy consumption, but they lack a strict quantitative model similar to PUE that would be convenient as a legislative control measure to express exactly how much they alter  $CO_2e$  emission levels. From the  $CO_2e$  management perspective, these methods work more in a best-effort manner, attempting first and foremost to satisfy SLA constrains.

In [16] the HGreen heuristic is proposed to schedule batch jobs on the greenest resource first, based on prior energy efficiency benchmarking of all the nodes, but not how to optimize a job once it is allocated to a node - how much of its resources is it allowed to consume. A similar multiple-node-oriented scheduling algorithm is presented in [22].

The work described in [15] has the most similarities with ours, since it also balances SLA and energy constraints and even describes energy consumption using a similar, linear model motivated by dynamic voltage scaling, but no consideration of  $CO_2e$  management was made inside the model.

A good overview of cloud computing and sustainability is given in [18], with explanations of where cloud computing stands in regard to  $CO_2e$  emissions. Green policies for scheduling are proposed that, if accepted by the user, could greatly increase the efficiency of cloud computing and reduce  $CO_2e$  emissions. Reducing emissions is not treated as a source of profit and a possible way to balance SLA violations, though, but more of a general guideline for running the data center to stay below a certain threshold.

## 2.2 Energy efficiency legislation

Measures of controlling the energy efficiency in data centers do exist – metrics such as power usage efficiency (PUE) [10], carbon usage efficiency (CUE), water usage efficiency (WUE) [9] and others have basically become the industry standards through the joint efforts of policy makers and cloud providers gathered behind The Green Grid consortium [8]. The problem with these metrics, though, is that they only focus on the infrastructure efficiency – turn as much energy as possible into computing inside the IT equipment. Once the power gets to the IT equipment, though, all formal energy efficiency regulation stops, making it more of a black-box approach. For this reason, an attempt is made in our work to bring energy efficiency control to the interior operation of cloud computing - resource scheduling. So far, the measurement and control of even such a basic metric as PUE is not mandatory. It is considered a best practice, though, and agencies such as the U.S. Environmental Protection Agency (EPA) encourage data centers to measure it by rewarding the best data centers with the Energy Star award [2].

## 3 Application of the Kyoto protocol to the clouds

The Kyoto protocol [19] commits involved countries to stabilize their greenhouse gas (GHG) emissions by adhering to the measures developed by the United Nations Framework Convention on Climate Change (UNFCCC) [12]. These measures are commonly known as *the cap-and-trade system*. It is based on setting national emission bounds - caps, and establishing international emission markets for trading emission surpluses and emission deficits. This is known as *certified emission reductions* or *credits for emission reduction* (CERs). This trade rewards countries which succeeded in reaching their goal with profits from selling CERs and allows those who did not to make up for it financially by buying CERs. The European Union Emission Trading System (EU ETS) is an example implementation of an emission trading market [4]. Through such markets, CERs converge towards a relatively constant market price, same as all the other tradable goods.

Individual countries control emissions among their own large polluters (individual companies such as power generation facilities, factories. . .) by distributing the available caps among them. In the current implementation, though, emission caps are only set for entities which produce more than 25 MtCO<sub>2</sub>e/year [3]. This excludes individual data centers which have a carbon footprint in the ktCO<sub>2</sub>e/year range [1].

It is highly possible, though, that the Kyoto protocol will expand to smaller entities such as cloud providers to cover a larger percentage of polluters and to increase the chance of global improvement. One such reason is that currently energy producers take most of the weight of the protocol as they cannot pass the responsibilities on to their clients (some of which are quite large, such as data centers). In 2009, three companies in the EU ETS with the largest shortage of carbon allowances were electricity producers [5]. Another indicator of the

justification of this forecast is that some cloud providers, such as Google already participate in emission trading markets to achieve carbon neutrality [7].

For this reason, we hypothesize in this paper that cloud providers are indeed part of an emission trading scheme and that  $CO_2e$  emissions have a market price.

## 4 Wastage-penalty balance in a Kyoto-compliant cloud

In this section we present our  $CO_2e$ -trading model that is to be integrated with cloud computing. We show how an economical balance can be found in it. Lastly, we give some discussion as to how such information might be integrated into a scheduler to make it more energy and cost efficient.


### 4.1 The $CO_2e$ -trading model

The goal of our model is to integrate the Kyoto protocol's  $CO_2e$  trading mechanism with the existing cloud computing service-oriented paradigm. At the same time we want to use these two aspects of cloud computing to express an economical balance function that can serve us to make better decisions in the scheduling process.

The model diagram in Fig. 4.1 shows the entities in our model and their relations. A cloud offers some computing resources as services to its users and they in turn pay the cloud provider for these services. Now, a cloud is basically some software running on machines in a data center. To operate, a data center uses electrical energy that is bought from an energy producer. The energy producers are polluters as they emit  $CO_2e$  into the atmosphere. As previously explained, to mitigate this pollution, energy producers are bound by the Kyoto protocol to keep their  $CO_2e$  emissions below a certain threshold and buy CERs for all the excess emissions from other entities that did not reach their caps yet over the emission trading market (ETM). This is illustrated by getting negative CERs (-CERs) for  $CO_2e$  responsibilities and having to buy the same amount of positive CERs (+CERs) over the ETM. It does not make any real difference for our model if an entity reaches its cap or not, as it can sell the remaining  $CO_2e$  allowance as CERs to someone else over the ETM.

Most importantly, this means that  *$CO_2e$  emissions an entity is responsible for have a price*. The other important thing to state in our model is that  *$CO_2e$  emission responsibilities for the energy that was bought is transferred from the energy producer to the cloud provider*. This is shown in Fig. 4.1 by energy producers passing some amount of -CERs to the cloud provider along with the energy that was bought. The cloud provider then has to buy the same amount of +CERs via the ETM (or he will be able to sell them if he does not surpass his cap making them equally valuable).

The consequences of introducing this model are that three prices influence the cloud provider: (1) energy cost; (2)  $CO_2e$  cost; (3) service cost. To maximize profit, the cloud provider is motivated to decrease energy and  $CO_2e$  costs and



**Fig. 1.** Cloud computing integrated with the Kyoto protocol's emission trading scheme

maximize earnings from selling his service. Since the former is achieved by minimizing resource usage to save energy and the latter by having enough resources to satisfy the users, they are conflicting constraints. Therefore, an economical balance is needed to find exactly how much resources to provide.

The service costs are much bigger than both of the other two combined (that is the current market state at least, otherwise cloud providers would not operate), so they cannot be directly compared, but there are different ways a service can be delivered, depending on how the cloud schedules resources. The aim of a profit-seeking cloud provider is to deliver just enough resources to the user so that his needs are fulfilled and that the energy wastage stays minimal. If a user happens to be tricked out of too much of the resources initially sold to him, a service violation occurs and the cloud provider has to pay a penalty price. This means that we are comparing the energy wastage price with the occasional violation penalty. This comparison is the core of our wastage-penalty model and we will now explain how can a wastage-penalty economical balance be found.

#### 4.2 Penalty-wastage resource balancing model

As was briefly sketched in the introduction, the main idea is to push cloud providers to follow the users' demands more closely, avoiding too much resource


over-provisioning, thus saving energy. We do this by introducing additional cost factors that the cloud provider has to pay if he wastes too much resources – the energy and  $CO_2e$  costs shown in Fig. 4.1, encouraging him to breach the agreed service agreements and only provide what is actually needed. Of course, the cloud provider will not breach the agreement too much, as that could cause too many violation detections (by a user demanding what cannot be provided at the moment) and causing penalty costs. We will now expand our model with some formal definitions in the cloud-user interface from Fig. 4.1 to be able to explicitly express the wastage-penalty balance in it.

We explain the trading model where we assume a situation with one cloud provider and one cloud user. The cloud provides the user with a single, abstract resource that constitute its service (it can be the amount of available data storage in GB, for example). To provide a certain amount of this resource to the user in a unit of time, a proportional amount of energy is consumed and indirectly a proportional amount of  $CO_2e$  is emitted. An example resource scheduling scenario is shown in Fig. 4.2. An SLA was signed that binds the cloud provider to provide the user a constant resource amount,  $r_{agreed}$ . The cloud provider was paid for this service in advance. A user uses different resource amounts over time. At any moment the  $R_{demand}$  value is what the user requires. To avoid over-provisioning the provider does not actually provision the promised resource amount all the time, but instead adapts this value dynamically,  $r_{provisioned}$  is the resource amount allocated to the user in a time unit. This can be seen in Fig. 4.2 as  $r_{provisioned}$  increases from  $t_1$  to  $t_2$  to adapt to a sudden rise in  $R_{demand}$ .

As we can not know how the user’s demand changes over time, we will think of  $R_{demand}$  as a random variable. To express  $R_{demand}$  in an explicit way, some statistical method would be required and research of users’ behaviour similar to that in [14] to gather real-life data regarding cloud computing resource demand. To stay on a high level of abstraction, though, we assume that it conforms to some statistical distribution and that we can calculate its mean  $\bar{R}_{demand}$  and its maximum  $max(R_{demand})$ . To use this solution in the real world, an appropriate distribution should be input (or better yet – one of several possible distributions should be chosen at runtime that corresponds to the current user or application profile). For most of the statistical distributions we know at least the random variable’s expected value  $E$  and variance  $V$  and we can express  $\bar{R}_{demand}$  as the expected value  $E(R)$  and  $max(R_{demand})$  as the sum of  $E(R) + V(R)$  with a limited error.

We denote the energy price to provision the whole  $r_{agreed}$  resource amount per time unit  $c_{en}$  and similarly the  $CO_2e$  price  $c_{co_2}$ . By only using the infrastructure to provision an amount that is estimated the user will require, not the whole amount, we save energy that would have otherwise been wasted and we denote this evaded wastage cost  $c_{wastage}$ . Since  $c_{wastage}$  is a fraction of  $c_{en} + c_{co_2}$ , we can use a percentage  $w$  to state what percentage is wasted:

$$c_{wastage} = w * (c_{en} + c_{co_2}) \quad (1)$$



**Fig. 2.** Changes in the *provisioned* and *demand* resource amounts over time

We know what the extreme cases for  $w$  should be – 0% for provisioning approximately what is needed,  $\bar{R}_{demand}$ ; and the percentage equivalent to the ratio of the distance between  $\bar{R}_{demand}$  and  $r_{agreed}$  to the total amount  $r_{agreed}$  if we provision  $r_{agreed}$ :

$$w = \begin{cases} 1 - \frac{\bar{R}_{demand}}{r_{agreed}}, & \text{if } r_{provisioned} = r_{agreed} \\ 0, & \text{if } r_{provisioned} = \bar{R}_{demand} \end{cases} \quad (2)$$

We model the distribution of  $w$  between these extreme values using linear interpolation:

$$w = \frac{r_{provisioned} - \bar{R}_{demand}}{r_{agreed}} \quad (3)$$

If we apply 3 to 1 we get:

$$C_{wastage} = \frac{r_{provisioned} - \bar{R}_{demand}}{r_{agreed}} * (c_{en} + c_{co2}) \quad (4)$$

If a user demands more resources than the provider has provisioned, an SLA violation occurs. The user gets only the provisioned amount of resources in this case and the provider has to pay the penalty cost  $C_{penal}$ . While  $c_{en}$  and  $c_{co2}$  can be considered constant for our needs,  $C_{penal}$  is a random variable, because it



depends on the user's behaviour which we can not predict with 100% accuracy, so we will be working with  $E(C_{penal})$ , its expected value.

Now that we have identified the individual costs, we can state our goal function. If the cloud provider provisions too much resources the  $c_{wastage}$  wastage cost is too high. If on the other hand he provisions too little resources, tightens the grip on the user too much, the  $E(C_{penal})$  penalty cost will be too high. The economical balance occurs when the penalty and wastage costs are equal - it is profitable for the cloud provider to breach the SLA only up to the point where penalty costs exceed wastage savings. We can express this economical balance with the following equation:

$$c_{wastage} = E(C_{penal}) + [\text{customer satisfaction factor}] \quad (5)$$

The *[customer satisfaction factor]* could be used to model how our promised-provisioned manipulations affect the user's happiness with the QoS and would be dependant of the service cost (because it might influence if the user would be willing to pay for it again in the future). For simplicity's sake we will say that this factor equals 0, getting:

$$c_{wastage} = E(C_{penal}) \quad (6)$$

$E(C_{penal})$ , the expected value of  $C_{penal}$  can be calculated as:

$$E(C_{penal}) = p_{viol} * c_{viol} \quad (7)$$

where  $c_{viol}$  is the constant cost of a single violation (although in reality probably not all kinds of violations would be priced the same) and  $p_{viol}$  is the probability of a violation occurring. This probability can be expressed as a function of  $r_{provisioned}$ ,  $r_{agreed}$  and  $R_{demand}$ , the random variable representing the user's behaviour:

$$p_{viol} = f(\bar{r}_{provisioned}, r_{promise}, R_{demand}) \quad (8)$$

Again, as for  $w$ , we know what extreme values for  $p_{viol}$  we want. If 0 is provisioned, we have 100% violations and if  $max(R_{demand})$  is provisioned, we have 0% violations:

$$p_{viol} = \begin{cases} 100\%, & \text{if } r_{provisioned} = 0 \\ 0\%, & \text{if } r_{provisioned} = max(R_{demand}) \end{cases} \quad (9)$$

and if we assume a linear distribution in between we get:

$$p_{viol} = 1 - \frac{r_{provisioned}}{max(R_{demand})} \quad (10)$$

Now, we can combine equations 4, 6, 7 and 10 to get a final expression for  $r_{provisioned}$ :

$$r_{provisioned} = \frac{\max(R_{demand}) * [(\bar{R}_{demand} * (c_{en} + c_{co2}) + r_{agreed} * c_{viol})]}{\max(R_{demand}) * (c_{en} + c_{co2}) + r_{agreed} * c_{viol}} \quad (11)$$

This formula is basically *the economical wastage-penalty balance*. Since all the parameters it depends on are constant (as long as the demand statistic stays the same) shows how much on average should a cloud provider breach the promised resource amounts when provisioning resources to users so that the statistically expected costs for SLA violation penalties do not surpass the gains from energy savings. Vice versa also holds - if a cloud provider provisions more resources than this wastage-penalty balance, he pays more for the energy wastage (energy and  $CO_2e$  price), than what he saves on SLA violations.

### 4.3 Heuristics for scheduling optimisation with integrated emission management

In this section we discuss a possible application of our penalty-wastage model for the implementation of a future-generation data center. Knowing the economical wastage-penalty balance, heuristic functions can be used to optimize resource allocation to maximize the cloud provider's profit by integrating both service and violation penalty prices and energy and  $CO_2e$  costs. This is useful, because it helps in the decision-making process when there are so many contradicting costs and constraints involved.

A heuristic might state: "try not to provision more than  $\pm x\%$  resources than the economical wastage-penalty balance". This heuristic could easily be integrated into existing scheduling algorithms, such as [21, 20] so that the cloud provider does not stray too much from the statistically profitable zone without deeper knowledge about resource demand profiles. The benefits of using our wastage-penalty model are:

- a new, expanded cost model covers all of the influences from Fig. 4.1
- $CO_2e$ -trading schema-readiness makes it easier to take part in emission trading
- a Kyoto-compliant scheduler module can be adapted for use in resource scheduling and allocation solutions
- the model is valid even without Kyoto-compliance by setting the  $CO_2e$  price  $c_{co2}$  to 0, meaning it can be used in traditional ways by weighing only the energy wastage costs against service violation penalties.

The wastage-penalty balance in 11 is a function of significant costs and the demand profile's statistical properties:

$$r_{provisioned} = g(\max(R_{demand}), \bar{R}_{demand}, r_{agreed}, c_{en}, c_{co2}, c_{viol}) \quad (12)$$

This function enables the input of various statistical models for user or application demand profiles ( $\max(R_{demand})$  and  $\bar{R}_{demand}$ ) and energy ( $c_{en}$ ),  $CO_2e$

( $c_{CO_2}$ ) and SLA violation market prices ( $c_{viol}$ ). With different input parameters, output results (e.g. energy savings, environmental impact, SLA violation frequency) can be compared. This would allow cloud providers and governing decision-makers to simulate the effects of different scenarios and measure the influence of individual parameters, helping them choose the right strategy.

## 5 Conclusion and Future Work

In this paper we presented a novel approach for Kyoto protocol-compliant modeling of data centers. Firstly, we presented a conceptual model for  $CO_2e$  trading compliant to the Kyoto protocol's emission trading scheme. Thereby, we consider *emission trading market (ETM)* where *credits for emission reduction (CERs)* are traded between data centers resulting in either in- or outsourcing of applications in order to achieve a  $CO_2e$  equilibrium. Finally, we proposed a *wastage - penalty* model that can be used as a basis for the implementation of Kyoto protocol-compliant scheduling and pricing models.

In the future we plan to implement prototype scheduling algorithms for the ETM considering self-adaptable Cloud infrastructures.

## Acknowledgements

The work described in this paper has been funded by the Vienna Science and Technology Fund (WWTF) through project ICT08-018 and by the TU Vienna funded HALEY project (Holistic Energy Efficient Management of Hybrid Clouds).

## References

1. Data center energy efficiency calculator, <http://www.42u.com/efficiency/energy-efficiency-calculator.htm>
2. ENERGY STAR data center energy efficiency initiatives : ENERGY STAR, [http://www.energystar.gov/index.cfm?c=prod\\_development.server\\_efficiency](http://www.energystar.gov/index.cfm?c=prod_development.server_efficiency)
3. Environment agency - phase III 2013 to 2020, [http://www.environment-agency.gov.uk/business/topics/pollution/113457.aspx#What\\_we\\_need\\_you\\_to\\_do\\_now](http://www.environment-agency.gov.uk/business/topics/pollution/113457.aspx#What_we_need_you_to_do_now)
4. EU emissions trading system - department of energy and climate change, [http://www.decc.gov.uk/en/content/cms/emissions/eu\\_ets/eu\\_ets.aspx](http://www.decc.gov.uk/en/content/cms/emissions/eu_ets/eu_ets.aspx)
5. EU ETS 2009 company ratings, <http://www.carbonmarketdata.com/cmd/publications/EU%20ETS%202009%20Company%20Rankings%20-%20June%202010.pdf>
6. Gartner estimates ICT industry accounts for 2 percent of global CO2 emissions, <http://www.gartner.com/it/page.jsp?id=503867>
7. Google's PPAs: what, how and why, [http://static.googleusercontent.com/external\\_content/untrusted\\_dlcp/www.google.com/en/us/green/pdfs/renewable-energy.pdf](http://static.googleusercontent.com/external_content/untrusted_dlcp/www.google.com/en/us/green/pdfs/renewable-energy.pdf)

8. The green grid, <http://www.thegreengrid.org/>
9. The green grid > CUEandWUE, <http://www.thegreengrid.org/Global/Content/TechnicalForumPresentation/2011TechForumCUEandWUE>
10. The green grid > the green grid metrics: Describing data center power efficiency, <http://www.thegreengrid.org/en/Global/Content/white-papers/Green-Grid-Metrics>
11. IPCC - intergovernmental panel on climate change, [http://www.ipcc.ch/publications\\_and\\_data/publications\\_ipcc\\_fourth\\_assessment\\_report\\_synthesis\\_report.htm](http://www.ipcc.ch/publications_and_data/publications_ipcc_fourth_assessment_report_synthesis_report.htm)
12. Kyoto protocol, [http://unfccc.int/kyoto\\_protocol/items/2830.php](http://unfccc.int/kyoto_protocol/items/2830.php)
13. Kyoto protocol to the united nations framework convention on climate change, [http://unfccc.int/essential\\_background/kyoto\\_protocol/items/1678.php](http://unfccc.int/essential_background/kyoto_protocol/items/1678.php)
14. Beauvisage, T.: Computer usage in daily life. In: Proceedings of the 27th international conference on Human factors in computing systems. p. 575–584. CHI '09, ACM, New York, NY, USA (2009)
15. Beloglazov, A., Buyya, R.: Adaptive threshold-based approach for energy-efficient consolidation of virtual machines in cloud data centers. In: Proceedings of the 8th International Workshop on Middleware for Grids, Clouds and e-Science. p. 4:1–4:6. MGC '10, ACM, New York, NY, USA (2010)
16. Coutinho, F., Carvalho, L.A.V.d., Santana, R.: A workflow scheduling algorithm for optimizing Energy-Efficient grid resources usage. In: 2011 IEEE Ninth International Conference on Dependable, Autonomic and Secure Computing (DASC). pp. 642–649. IEEE (Dec 2011)
17. Ellerman, D., Joskow, P.: The european union's emissions trading system in perspective, prepared for the pew center on global climate change, <http://www.c2es.org/eu-ets>
18. Garg, S.K., Yeo, C.S., 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 (Jun 2011)
19. Grubb, M., Vrolijk, C., Brack, D.: *Kyoto Protocol: A Guide and Assessment*. Earthscan Ltd, London (Jul 1999)
20. Maurer, M., Brandic, I., Sakellariou, R.: Simulating autonomic SLA enactment in clouds using case based reasoning. *Towards a Service-Based Internet* p. 25–36 (2010)
21. Maurer, M., Brandic, I., Sakellariou, R.: Enacting SLAs in clouds using rules. *Euro-Par 2011 Parallel Processing* p. 455–466 (2011)
22. Wang, Z., Zhang, Y.: Energy-Efficient task scheduling algorithms with human intelligence based task shuffling and task relocation. pp. 38–43. IEEE (Aug 2011)