Quantitative Analysis of the Effects of Different Carbon Tax Levels on Emissions and Costs of Data Centers

Sascha Bosse, Abdulrahman Nahhas and Klaus Turowski

Magdeburg Research and Competence Cluster for Very Large Business Applications
Otto-von-Guericke-University Magdeburg, Universitaetsplatz 2, 39106 Magdeburg, Germany
{sascha.bosse, abdulrahman.nahhas, klaus.turowski}@ovgu.de

Abstract. Emissions of greenhouse gases (GHG) have to be reduced to limit the impacts of climate change. For that reason, the introduction of carbon taxes has been discussed or performed in many countries. Data centers are accounting for an increasing fraction of GHG emissions, so that carbon taxation may lead to reduced emissions. In this paper, the effect of different carbon tax levels is analyzed in experiments based on real-world workload from 20 data centers hosting enterprise systems. From the results, it can be concluded that optimization potential can be addressed with server consolidation, limiting the additional costs to be expected. Additionally, the used power mix and the depreciation period have a strong influence on the additional cost as well as the optimization potential regarding emissions.

Keywords: Carbon tax, Data center management, Server consolidation, Greenhouse gas emissions

1 Introduction

Climate change is undoubtedly the greatest challenge of our time. Severe effects on ecology and economy are to be expected due to, for example, rising sea levels or extreme weather conditions [1]. As of today, the Paris Agreement of the United Nations Framework Convention on Climate Change, which aims at limiting the increase in average global temperature to 1.5 Kelvin [2], has been ratified by 187 nations[1]. However, the average global temperature in the last 30 years has already increased by 1.0 Kelvin in comparison with pre-industrial levels. Thus, anthropogenic greenhouse gas (GHG) emissions, measured in carbon dioxide equivalent (CO2eq), have to be reduced to zero until 2050 to achieve this goal very likely [1]. The federal government of Germany, for instance, wants to reduce GHG emissions as early as possible by more than 13 percent on the way to achieve GHG neutrality in 2050 [3].

The major problem when reducing GHG emissions in a capitalist system is that the price of any product does not really contain its environmental cost, e.g. due to subsidiaries or outsourcing [4, 5]. While regulations could be one way to reduce GHG

[1] https://treaties.un.org/Pages/ViewDetails.aspx?src=TREATY&mtdsg_no=XXVII-7-d&chapter=27&clang=_en

15th International Conference on Wirtschaftsinformatik,
March 08-11, 2020, Potsdam, Germany

https://doi.org/10.30844/wi_2020_m2-bosse

emissions, experts argue to utilize economic mechanisms such as trading systems or carbon taxes [1, 6]. In that context, the European Emission Trading System (ETS) has been established in 2005 to reduce GHG emissions [7]. However, the ETS has been criticized, e.g. for free certificates, missing trader frauds, or price volatility [7, 8] and the amount of GHG emissions in the EU-28 remained constant in the last years (and is still increasing worldwide) [9]. For that reasons, carbon taxes are nowadays regarded as an effective, additional instrument to reduce emissions [8, 10, 11] and have been implemented by 46 countries in 2019 ranging from 0.10 € to 120 € per ton CO2eq [11]. In other countries, such as Germany or China, a carbon tax is currently being discussed.

One of the industrial sector with the highest increase in GHG emissions is the ICT sector [12]. The crypto currency boom alone increased the energy consumption of data centers by more than 43 TWh per year [13]. Also the enterprise system sector, as the IT sector with highest growth rates [14], can contribute effectively to the reduction of GHG emissions.

In the context of server consolidation, the costs of a data center are optimized by co-allocating orthogonal workload patterns to minimize the number of required computing resources [15, 16]. Live migration introduces even more flexibility in server consolidation [15], but cannot always be applied, especially for enterprise systems [17].

Since servers can differ strongly in their emissions and power consumption behavior [18] a carbon tax will likely affect these efforts. As a result, low-emissive allocations should be preferred. Therefore, this paper aims at answering the following research question: *which effects are different carbon tax levels inducing to the costs and emissions of data centers that apply server consolidation?* For that reason, computational experiments are performed in which server consolidation problems are solved based on real-world workload from 20 data centers hosting enterprise systems. In these experiments, five tax levels are considered (0-180 € per ton CO2eq) as well as two power sources (fossil and renewable) and two depreciation periods (3 and 5 years).

2 Related Work

Although the effect of carbon taxation to data centers using server consolidation has not been analyzed before, there are several studies available on a macroeconomic level. These works often use simulation to determine the effects of a carbon tax which is parameterized with empirical data, e.g. [19–22]. Since carbon taxes are implemented in more and more countries, also pure empirical analyses are available, both on the macro- and microeconomic level, e.g. [23, 24]. For predicting microeconomic effects, however, domain-specific models as in [25] have to be used. Therefore, the server consolidation problem itself must be subject to investigations about carbon taxes, as these may affect the objective function directly.

Although most server consolidation approaches derive from the bin packing optimization problem [26], these can differ in their results due to their objective

https://doi.org/10.30844/wi_2020_m2-bosse

function. While traditional approaches for server consolidation focused on the minimization of cost [27], the majority of recent approaches minimizes the energy consumption of the data center [15].

Although consumed energy is an important cost factor for data centers [28], the emissions accounted for the energy source are not the only source of emissions in its lifecycle [29]. Next to data center construction, this refers mainly to emissions caused by manufacturing, packaging, transportation, storage, and disposal of used IT components [30, 31]. While some server manufacturers have published lifecycle emission analyses, these are not completely reliable due to variance in usage profile [32]. Furthermore, the increasing energy efficiency of servers [30] leads to the fact that the fraction of emissions caused by usage is decreasing. Therefore, optimization efforts targeting at the minimization of costs in the presence of a carbon tax cannot rely solely on the minimization of energy consumption. Instead, these should be integrated into cost minimization approaches that are also considering the acquisition of IT components.

3 Experimental Design

The server consolidation problems which are to be solved in the experiments are derived from Speitkamp and Bichler [16]. This problem formulation is suitable for enterprise systems that are not hardware-virtualized to avoid performance drawbacks of the virtualization layer [33]. Seasonal workload patterns of services can be utilized to consolidate servers to minimize total cost without performance degradations, which should be avoided especially in a business context [34]. Workload profiles $W = (w_1, \ldots)$ for 20 real-world cases hosting business applications have been provided by an industry partner.

However, in order to illustrate the effects of a carbon tax to these cases, the following assumptions are made. First, since power usage profiles in combination with carbon footprint data are only available for a few server types, only these have been used in the experiments. Second, it is assumed that a carbon tax is implemented independently of a trading mechanism and is completely passed through by power suppliers and server manufacturers. Third, costs and emissions accounting for data center construction are not considered as these are hard to estimate. Finally, it is assumed that all parameters stay constant for the whole depreciation period, so that neither workloads, power mix parameters nor the allocation change.

The required parameters and functions to define the server consolidation problem are presented in Table 1 as well as in the following.

The optimization problem is solved in different scenarios, depending on the tax level tax, the depreciation period y as well as the used power mix described by (p, e). Here, p is the price per kWh in € and e denotes the emissions per kWh in kgCO2eq. The workload for a service $w_i = \big((c_1, m_1) \ldots (c_{24}, m_{24})\big)$ is characterized by its computing and memory demand over 24 hours of the day, forming a seasonal workload profile. Computing demands and capacities are given in the SAP

https://doi.org/10.30844/wi_2020_m2-bosse

Application Performance Standard (SAPS) which is an application-dependent, but hardware-independent measure for computing resources [35].

Table 1. Used symbols and description

Symbol	*Description*
W	The set of workload profiles
c_t^w	Computing demand in SAPS of workload w at hour $t, 1 \leq t \leq 24$
m_t^w	Memory demand in MB of workload w at hour $t, 1 \leq t \leq 24$
y	Length of depreciation period in years
p	Price for energy per kWh in €
e	Emissions for energy per kWh in kgCO2eq
a	An allocation of workloads to servers
tax	Carbon tax level in €/tCO2eq
PUE	Power-usage-effectiveness factor
S	The set of servers to which workloads are allocated
$TCY(a)$	Total costs per year for an allocation
$acCost(a)$	Acquisition costs for an allocation
$acCO2(a)$	Non-usage emissions for an allocation
$opCost(a)$	Operational costs for an allocation
$opCO2(a)$	Usage emissions for an allocation
$E_y(a)$	Energy consumption of an allocation for a period y
$util(s, a, t)$	Utilization of a server s at hour t for an allocation
$power_s(u)$	Power consumption of a server s at utilization level u
cap_s	Computing capacity of server s in SAPS

Services should then be allocated to servers so that their workload over the day must not exceed server capacity. The available servers are created using five server types which are presented in Table 2 and in Figure 1. In the latter, the power consumption functions $power_s$ are compared depending on the computing load level. It can be stated that the server types differ significantly in their power consumption behavior as well as in their price and non-usage emissions (w.r.t. resource capacity). Thus, type A and B represent expensive, but energy-efficient server types while types C, D, and E represent cheap alternatives with a more inefficient energy consumption.

Table 2. Server types used in the experiments

Label	*Introduction year*	*Price in €*	*Non-usage emissions in kgCO2eq*	*Computing capacity in SAPS*	*Memory capacity in GB*
A	2018	899.00	219	5,504	8
B	2017	4,599.00	500	11,543	16
C	2011	293.00	400	4,017	24
D	2008	235.50	370	3,868	8
E	2010	173.95	543	6,853	4

https://doi.org/10.30844/wi_2020_m2-bosse

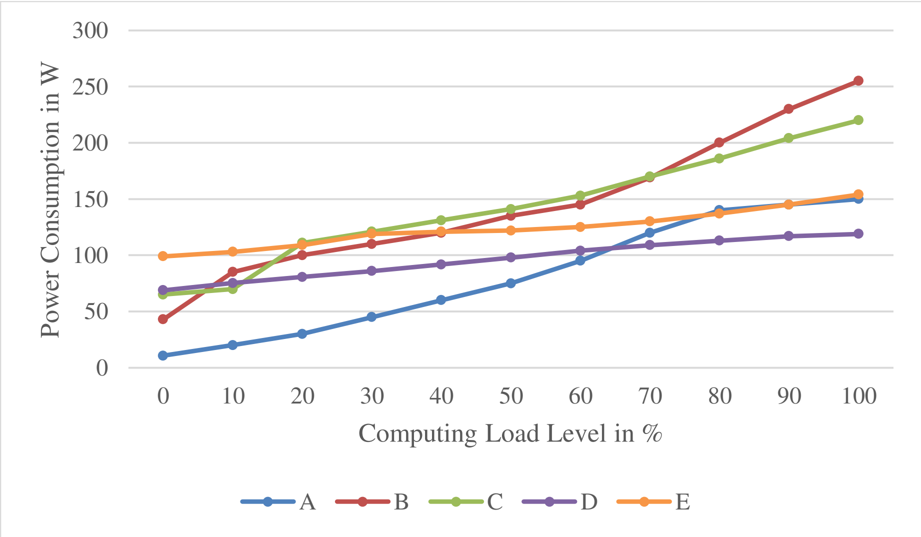

Figure 1. Power consumption profiles for used server types[2]

The objective function to be minimized in the optimization problem is formulated in Equation 1 (total cost per year), depending on the allocation of services to servers a:

$$TCY(a) = \frac{1}{y} \cdot (acCost(a) + acCO2(a) \cdot tax + opCost(a) + opCO2(a) \cdot tax) \quad (1)$$

In this function, the acquisition costs are the sum of acquisition costs for each used server, similar to the computation of acquisition emissions. For operational costs and emissions, the energy consumption of the allocation is computed according to Equation 2, depending on the power-usage effectiveness PUE as well as server computing utilization level. The latter is in turn computed according to Equation 3.

$$E_y(a) = PUE \cdot 365y \cdot \sum_{t=1}^{24} \sum_{s \in S} power_s\big(util(s,a,t)\big) \; / \; 1000 \quad (2)$$

$$util(s,a,t) = \big(\sum_{w \in W, a(w)=s} c_t^w\big)/cap_s \quad (3)$$

Operational cost resp. emissions are then the product of the energy consumption and kWh price resp. emissions. As a constraint, an allocation must not lead to a utilization above 100%.

The PUE is set to 2.0, which is reported as average [30], and for the depreciation period y, three and five years are used in the experiments. In order to demonstrate the behavior in edge cases, two extremely different power mixes have been defined according to German supplier information[3]. These are compared in Table 3. The fossil

---

[2] https://www.spec.org/power_ssj2008/results/power_ssj2008.html
[3] https://web-api.vattenfall.de/service-apis/download/document/90b85850-88ab-4275-9b18-3423693f68e5

https://doi.org/10.30844/wi_2020_m2-bosse

mix produces high emissions due to combustion of coal and gas while the renewable mix is considered with zero emissions.

Table 3. Power mixes used in the experiments

Label	*Price in €/kWh*	*Emissions in kgCO2eq/kWh*
Fossil	0.17090	0.6901
Renewable	0.23242	0.0000

In order to investigate the effect of different carbon tax scenarios, five different tax levels are chosen:

- a scenario with no carbon tax ($tax = 0$),
- a moderate carbon tax of 22 € as in France [11],
- 62 € as recommended by the International Monetary Fund [6],
- 120 € as in Sweden (highest carbon tax in OECD countries) [11]), and
- 180 € per ton CO2eq, which are the officially estimated climate costs in 2016 for Germany [36].

With an increasing tax level, the emissions from the used IT components increase the total costs of the data center. Thus, server consolidation will have a drift to prefer low-emissive allocations that may have higher costs without taxation but lead to lower costs in taxation scenarios. With an increasing tax level, the role of the consumed energy to the total costs will depend more and more on the usage emissions of the used power mix.

As exact methods may lead to execution problems for the large-sized cases due to scalability issues [37, 38], different heuristic (first- and best-fit-decreasing), meta-heuristic optimization approaches ((grouping) genetic algorithms), and hybrid strategies (genetic algorithm for first-/best-fit allocation) are applied for solving the generated server consolidation problems and the best feasible results are reported. Cf. [35, 38–40] for details regarding the solution algorithms.

4 Experimental Results

4.1 Results for All Cases

The results of all cases are presented in comparison to the optimization result of the baseline scenario. In this scenario, fossil power is used, and no carbon tax must be paid. In Table 4, the results of the baseline scenario are compared against the scenario with a carbon tax of 180 € in a 3-year depreciation period. The results of the 5-year depreciation period are presented in Table 5. For each case, the number of services to be allocated ($|W|$) and the cost per year for the baseline scenario are displayed. Baseline increase refers to the additional cost of the 180 € tax level if the same allocation would have been used. The column optimization potential indicates the fraction of additional cost that can be avoided due to optimization efforts. Similarly,

https://doi.org/10.30844/wi_2020_m2-bosse

the column GHG reduction states the relative reduction in emissions for the optimized allocation. The last three columns compare the baseline scenario to the scenario with renewable power and 180 € tax level. The cases, in which optimization potential has been addressed, are highlighted in gray.

Table 4. Results comparing baseline and 180€ tax scenarios for a 3-years depreciation period

#	$\|W\|$	*Baseline costs in € per year*	*Baseline increase in %*	*Opt. pot. in %*	*GHG reduct. in %*	*Baseline increase (ren.)*	*Opt. pot. (ren.)*	*GHG reduct. (ren.)*
1	5	441.79	62.1	4.0	37.0	33.5	0.0	0.0
2	10	705.03	62.8	22.5	53.9	34.3	0.0	0.0
3	9	785.08	62.3	12.6	25.4	33.8	1.4	58.3
4	13	827.65	63.9	23.0	69.0	40.9	2.6	5.9
5	11	931.98	66.6	0.0	0.0	36.1	0.0	0.0
6	15	1,215.43	62.1	7.5	16.2	33.7	0.0	0.0
7	10	1,339.42	64.0	0.0	0.0	34.3	0.0	0.0
8	27	1,429.30	62.9	19.2	51.8	34.3	0.0	0.0
9	80	2,203.53	63.7	0.0	0.0	34.2	0.0	0.0
10	36	2,218.52	56.1	0.0	0.0	30.2	0.0	0.0
11	28	2,224.91	57.4	21.7	40.1	31.8	8.3	55.2
12	42	2,433.79	61.7	4.6	28.9	33.8	0.0	0.0
13	18	2,682.93	29.8	5.2	15.6	17.0	1.1	43.6
14	80	2,956.29	27.5	0.0	0.0	14.8	0.0	0.0
15	30	5,215.68	47.3	6.9	7.2	25.5	6.9	10.0
16	28	6,328.91	34.7	0.0	0.0	18.6	0.0	0.0
17	27	7,610.76	13.0	0.0	0.0	7.3	0.0	0.0
18	16	10,476.44	27.3	4.9	9.0	14.7	3.6	3.6
19	44	12,271.32	40.8	34.2	46.6	22.1	24.8	43.0
20	51	32,923.89	18.9	0.0	0.0	10.3	0.0	0.0
∅	29	4,861.13	49.2	8.3	20.0	27.1	2.4	11.0

Due to higher acquisition and operation cost in the 180 € tax scenario, the maximum increase in cost per year is varying between 7 and 69% in relation to the baseline scenario. The additional costs to be expected depend heavily on the used power source and its associated emissions: for renewable power, the cost increase varies between 7 and 41 % (for fossil power 13 to 69%).

Nonetheless, in 16 of 20 cases, optimization potential can be addressed by changing the allocation to achieve lower costs in comparison to the baseline allocation. If the optimized allocation changes, the additional costs can be decreased by 0.5 to 46% (5% on average for all cases and scenarios). If the allocation changes, also significant emission savings can be achieved, starting by 4% to reaching up to 73% of original emissions (12.5% on average for all cases and scenarios).

The low or non-existing optimization potential in some cases can be explained by the fact that either a low-emissive allocation is preferred even without taxation or that an allocation with lower emissions would only be considered under tax levels higher than 180 €. Remarkably, emissions even increase in case 16 if a renewable power source is used in comparison to the baseline allocation. This can be explained by the fact that the optimization in the baseline scenario with fossil power puts more emphasis on usage costs. In the scenario of 180 € tax and renewable power, the acquisition costs account for a larger part of the yearly cost. Therefore, the allocation with higher emissions still leads to lower cost if tax is not higher.

Table 5. Results comparing baseline and 180€ tax scenarios for a 5-years depreciation period

#	$\|W\|$	*Baseline costs in € per year*	*Baseline increase in %*	*Opt. pot. in %*	*GHG reduct. in %*	*Baseline increase (ren.)*	*Opt. pot. (ren.)*	*GHG reduct. (ren.)*
1	5	389.21	39.7	0.0	0.0	23.2	0.0	0.0
2	10	604.32	32.1	0.0	0.0	17.2	0.0	0.0
3	9	688.70	50.8	0.0	0.0	26.8	0.0	0.0
4	13	673.85	51.1	46.2	56.9	31.7	42.8	72.5
5	11	869.72	68.8	0.0	0.0	36.0	0.0	0.0
6	15	1,082.86	56.1	0.0	0.0	29.5	0.0	0.0
7	10	1,237.55	67.0	13.5	19.9	34.9	7.8	26.5
8	27	1,243.15	33.2	0.0	0.0	17.7	0.0	0.0
9	80	2,031.20	66.8	17.1	31.0	34.8	4.6	36.0
10	36	1,957.72	61.4	0.0	0.0	32.0	0.0	0.0
11	28	1,860.93	39.7	0.0	0.0	20.7	0.0	0.0
12	42	2,162.59	45.9	4.4	34.2	25.4	0.0	0.0
13	18	1,984.27	37.9	12.6	13.9	20.6	14.1	43.6
14	80	2,184.16	36.0	4.0	26.8	18.8	0.0	0.0
15	30	4,286.56	51.6	4.8	7.5	26.9	3.3	4.7
16	28	4,907.08	43.3	8.7	15.7	22.6	0.5	-20.7
17	27	5,038.09	18.6	0.0	0.0	10.0	0.0	0.0
18	16	7,660.14	34.1	3.8	3.9	17.9	3.7	3.8
19	44	9,069.82	28.4	0.0	0.0	14.9	0.0	0.0
20	51	22,852.43	26.2	0.0	0.0	13.8	0.0	0.0
∅	29	3,639.22	44.4	5.8	10.5	23.8	3.8	8.3

However, the results show differences with respect to the depreciation period as well as to the used power mix. In Figure 2, the average savings in emissions depending on scenario and tax level are displayed with its 95%-confidence interval in blue. With a 5-year depreciation period and regenerative power, increasing tax levels have not led to savings in emissions for any case. However, as the red curve indicates, the allocation with lowest emissions has been used in 17 cases even without carbon taxation. To optimize the remaining three cases in terms of emissions, higher tax levels would be required. Also, for a three-year depreciation period and regenerative power, only in a few cases the allocations change depending on the tax level, leading to significant savings only for a 180 € tax level. However, in the scenarios with fossil power, emission savings can be achieved in many cases. The number of cases with the lowest emissions increases proportional to tax level, up to achieving allocations with lowest emissions in all cases for a 5-year depreciation period.

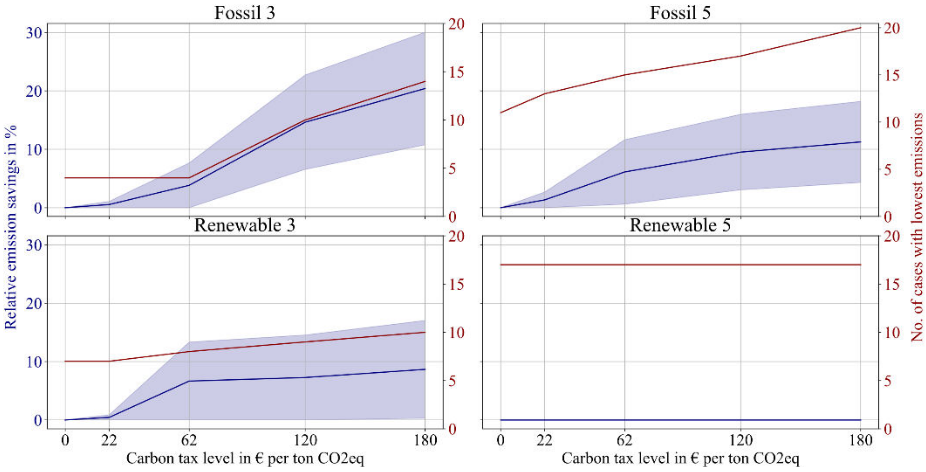

Figure 2. Relative emission savings with 95%-confidence interval (blue) and number of cases with lowest emissions (red) depending on tax level.

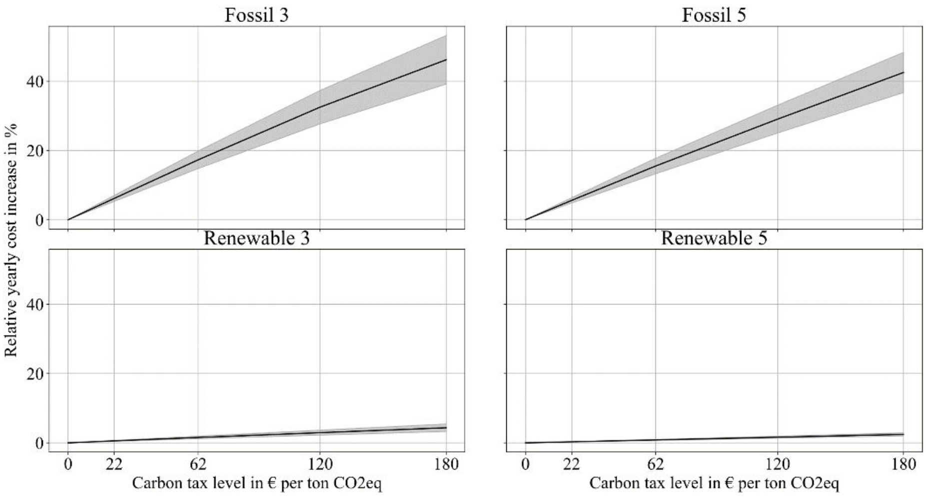

Figure 3. Increase in total yearly costs with 95%-confidence interval depending on tax level.

Why a carbon tax is more effective if high-emissive power sources are used can also be seen in Figure 3, in which the relative cost increase is displayed depending on the tax level. While costs will increase by 36%-53% on average in the scenarios with fossil power, the cost increase in the scenarios with regenerative power is limited to 2%-5% on average. Since server consolidation results adapt to the rising tax level, the relation between cost increase and tax level is not linear.

https://doi.org/10.30844/wi_2020_m2-bosse

4.2 Detailed Analysis of Two Cases

In order to get more insights regarding the optimization under different tax levels, two of the 20 cases presented above are analyzed in more detail. These are the cases 4 and 15 since these show significant optimization potential in some scenarios and represent different problem sizes.

In Fig. 2, the optimization results of case 4 are displayed for the different power mixes, depreciation periods, and tax levels. While the colored lines represent the objective values (yearly costs), the dashed lines compare these optimized values to the ones if the allocation from the zero tax scenario would be used. Next to each data point, the used server types in the allocation are presented ("AE^4" refers to one server of type A and four of type E are being used for the allocation).

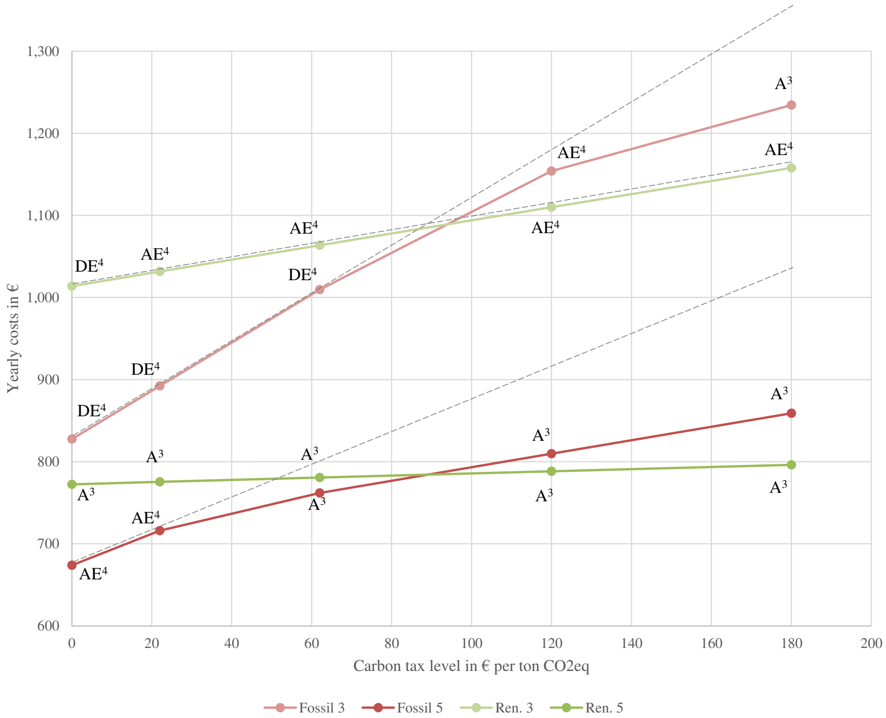

Figure 4. Optimization results for case 4

Comparing the dashed with the optimized curves, it becomes obvious that most optimization potential in this case can be achieved in the scenarios with fossil power. The allocation differs from the baseline allocation first at 120 € tax for a 3-year and at 62 € for a 5-year depreciation period. Additionally, the shift from fossil to renewable power becomes economically appropriate at a tax level of 120 € for both depreciation periods. In case a renewable power source is used in the beginning, the optimization potential will be significantly lower. Comparing the three different allocations that are reported, it can be stated that allocation "A^3" is the most expensive in acquisition

https://doi.org/10.30844/wi_2020_m2-bosse

costs but has lower acquisition emissions and power consumption than the other allocations. The allocation “DE^4” is the cheapest in terms of acquisition costs but has a high power consumption which is penalized in higher tax level scenarios, especially if fossil power is used. The allocation “AE^4” is a compromise of both extreme allocations and is used in scenarios in which the pressure of operational costs is not strong enough to justify the expensive allocation “A^3”.

Fig. 4 displays the results of case 15. Depending on the scenario, four different allocations are to be preferred from the cheapest and most power-consuming to the most expensive and least power-consuming. The allocation “BC^4DE^2” is recommended only in the baseline scenario. With higher tax levels, the allocation “ABC^3D^2E” is reported as optimal for a 3-year period. This also applies to the scenario with a 5-year period, zero tax, and fossil power. The allocation “A^2BC^3DE” is obtained in most of the other 5-year period scenarios, except for the 180 € tax and fossil power scenario, in which the allocation “$A^3BC^3E^2$” is recommended.

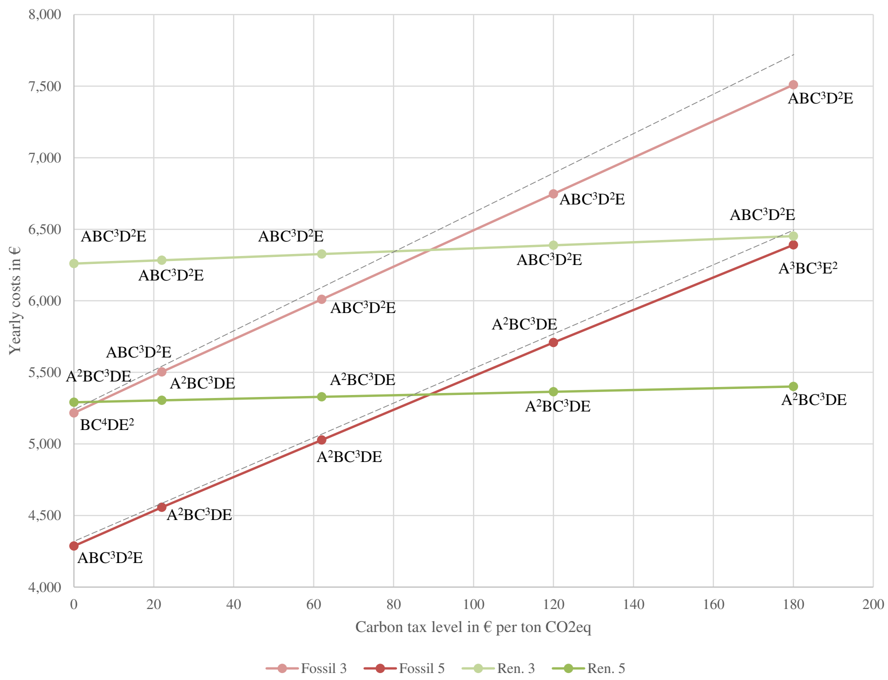

Figure 5. Optimization results for case 15

When the specific results of the analyzed cases are compared to the server types, it can be stated that the allocations mostly differ in the use of the server type A versus the types C, D, or E. While the latter ones represent rather outdated hardware with relatively high power consumption, but low acquisition cost, server type A has been recently introduced. Thus, it is more energy-efficient, but associated with higher acquisition costs. With an increasing tax level, the lower emissions of allocations relying on type A are more and more preferred. The same can be stated for the longer

https://doi.org/10.30844/wi_2020_m2-bosse

depreciation period, in which the acquisition costs per year are lower. However, this does not necessarily hold for the scenarios with a renewable power source as the lack of usage emissions limit the cost increase massively in comparison to the scenarios with a fossil power source.

Taking the results of all cases into account, the introduction of a carbon tax is most effective (from a GHG reduction point of view) the more emitting the used power source and the shorter the depreciation period. If the power source is low-emissive, a carbon tax will only affect the acquisition costs which considerably limits its influence on optimization efforts. With longer depreciation periods, operational costs are more dominant in the objective function even without a carbon tax. In both cases, consolidation approaches minimizing the power consumption will likely achieve the same results as cost minimization subject to carbon tax.

On the other hand, a carbon tax will likely drive data center managers to prefer longer depreciation periods and low-emissive power sources, since these instruments will reduce the additional cost. In all cases, the renewable power source would be preferred from an economic point of view if a carbon tax between 80 and 100 € is introduced. However, investments in expensive, but energy-efficient hardware are not always beneficial with rising carbon tax levels, especially if the used power source is low-emissive. This can be seen by the fact that the allocation with the least GHG emissions is not preferred in every case for the scenarios with a renewable power source. In these cases, the higher acquisition costs of the most energy-efficient allocation cannot be compensated even with increasing tax level.

It can be concluded that the introduction of a carbon tax will lead to increasing cost for data center providers. Server consolidation efforts can reduce these additional costs. In addition, considering longer depreciation periods or low-emissive power sources will reduce additional cost, especially for relatively high carbon tax levels. Combining these efforts, the carbon tax will have the desired ecological effect, driving data center providers to minimize GHG emissions.

5 Conclusion

Carbon taxes are regarded as an effective method to utilize economic mechanisms with the aim to reduce greenhouse gas emissions. While carbon taxes have already been introduced in 46 countries in different levels, in other countries it is heavily discussed. Data centers, as the backbone of the ongoing digitization, are emitting more and more greenhouse gases and can effectively contribute to reduce emissions by applying server consolidation in a low-emissive manner.

In this paper, the effect of different carbon tax levels to data centers has been analyzed. For that reason, workloads from 20 data centers hosting enterprise systems are used to define experiments. For the allocation of these workloads, five server types have been defined for which emission information was made available. In the experiments, five tax levels (0, 22, 62, 120, and 180 € per ton CO2eq), two extreme power mixes (fossil and renewable), and two depreciation periods (3 and 5 years) have been considered.

https://doi.org/10.30844/wi_2020_m2-bosse

While the results of the experiments are subject to some limitations (cf. Section 3), interesting conclusions have been drawn: first, a carbon tax can increase data center costs heavily depending on the used power source. While a carbon tax will likely increase server acquisition prices, the increase in energy price is proportional to the emissions per kWh of the power mix. Thus, a relevant carbon tax (above 62 €) will drive data centers to use low-emissive power sources. Furthermore, the experiments revealed that longer depreciation periods are to be preferred in order to save emissions. However, longer depreciation periods would also lead to higher uncertainty about future workload demands, maybe leading to more unused computing capacity and increasing emissions.

Second, significant optimization potential for the yearly costs can be addressed in some cases using server consolidation. These efforts will also lead to significant emission savings, especially if a high-emissive power source is used. Two exemplary cases have been presented to demonstrate the influence of the experimental parameters on the optimal allocation of workloads to servers. However, the existing variety in available server types could not be considered in this work due to lack of data, especially for non-usage emissions. This increases the complexity of the server consolidation problem, so that even more optimization potential may be addressed.

In order to provide more evidence for these conclusions, future work should focus on empirical investigations of the IT sector in countries with carbon taxation. This would also answer if the conclusions are transferable to other IT domains without seasonal workload pattern as it can be found in enterprise systems. Nevertheless, it can be stated that the introduction of a relevant carbon tax will likely have the desired effect in many cases, driving decision-makers to reduce emissions.

References

1. M. R. Allen, et al.: Global warming of 1.5°C. An IPCC Special Report on the impacts of global warming of 1.5°C above pre-industrial levels and related global greenhouse gas emission pathways - Technical Summary. Intergovernmental Panel on Climate Change (2018)
2. UN Framework Convention on Climate Change. United Nations, available at http://unfccc.int/resource/docs/2015/cop21/eng/l09r01.pdf (2015) (Accessed: 29.11.2019)
3. Klimaschutzziele Deutschlands. Umweltbundesamt der Bundesrepublik Deutschland, available at https://www.umweltbundesamt.de/daten/klima/klimaschutzziele-deutschlands (2019) (Accessed: 29.11.2019)
4. Bertram, C., Luderer, G., Pietzcker, R.C., Schmid, E., Kriegler, E., Edenhofer, O.: Complementing carbon prices with technology policies to keep climate targets within reach. Nature Climate Change 5, 235–239 (2015)
5. Okubo, T., Cole, M.A., Elliott, R.J.R.: Environmental Outsourcing. Research Institute of Economy, Trade and Industry Discussion Paper Series (2010)
6. Fiscal Policies for Paris Climate Strategies - From Principle to Practice. International Monetary Fund (2019)

https://doi.org/10.30844/wi_2020_m2-bosse

7. Gilbertson, T., Reyes, O., Lohmann, L.: Carbon trading: how it works and why it fails. Dag Hammarskjöld Foundation Uppsala (2009)
8. Stram, B.N.: A new strategic plan for a carbon tax. Energy Policy 73, 519–523 (2014)
9. Le Quéré, C., Andrew, R.M., Friedlingstein, P., others: Global carbon budget 2018. Earth System Science Data 10 (2018)
10. Ekins, P., Andersen, M.S., Vos, H.: Environmental taxes: Implementation and environmental effectiveness. Publications Office of the European Union (1996)
11. Group, W.B.: Carbon Pricing Watch 2019. World Bank Group, available at https://openknowledge.worldbank.org/handle/10986/31755 (2016) (Accessed: 29.11.2019)
12. Belkhir, L., Elmeligi, A.: Assessing ICT global emissions footprint: Trends to 2040 & recommendations. Journal of Cleaner Production 177, 448–463 (2018)
13. Kosharnaya, Y., Yanchenko, S., Kulikov, A.: Specifics of Data Mining Facilities as Energy Consumers. In: 2018 Dynamics of Systems, Mechanisms and Machines (Dynamics). 1–4. IEEE (2018)
14. Gartner: Worldwide IT Spending Forecast, available at https://www.gartner.com/en/newsroom/press-releases/2019-10-07-gartner-says-global-it-spending-to-grow-06-in-2019 (2019) (Accessed: 29.11.2019)
15. Lopez-Pires, F., Baran, B.: Virtual machine placement literature review. arXiv preprint arXiv:1506.01509. (2015)
16. Speitkamp, B., Bichler, M.: A Mathematical Programming Approach for Server Consolidation Problems in Virtualized Data Centers. IEEE Transactions on Services Computing 3, 266–278 (2010)
17. Rolia, J., Andrzejak, A., Arlitt, M.: Automating Enterprise Application Placement in Resource Utilities. In: Brunner, M. and Keller, A. (eds.) 14th IFIP/IEEE International Workshop on Distributed Systems: Operations and Management (DSOM). 118–129. Springer (2003)
18. Splieth, M., Bosse, S., Schulz, C., Turowski, K.: Analyzing the Effects of Load Distribution Algorithms on Energy Consumption of Servers in Cloud Data Centers. In: Proceedings of the 12th International Conference on Wirtschaftsinformatik (WI) (2015)
19. Bruvoll, A., Larsen, B.M.: Greenhouse gas emissions in Norway: do carbon taxes work? Energy policy 32, 493–505 (2004)
20. Davis, L.W., Kilian, L.: Estimating the effect of a gasoline tax on carbon emissions. Journal of Applied Econometrics 26, 1187–1214 (2011)
21. Lin, B., Li, X.: The effect of carbon tax on per capita CO2 emissions. Energy policy 39, 5137–5146 (2011)
22. Zhou, S., Shi, M., Li, N., Yuan, Y.: Impacts of carbon tax policy on CO2 mitigation and economic growth in China. Advances in Climate Change Research 2, 124–133 (2011)
23. Martin, R., De Preux, L.B., Wagner, U.J.: The impact of a carbon tax on manufacturing: Evidence from microdata. Journal of Public Economics 117, 1–14 (2014)
24. Zhao, Y.-H.: The study of effect of carbon tax on the international competitiveness of energy-intensive industries: an empirical analysis of OECD 21 countries, 1992-2008. Energy Procedia 5, 1291–1302 (2011)

https://doi.org/10.30844/wi_2020_m2-bosse

25. Brännlund, R., Nordström, J.: Carbon tax simulations using a household demand model. European Economic Review 48, 211–233 (2004)
26. Garey, M., Johnson, D.: Computers and Intractability: A Guide to the Theory of NP-Completeness. W.H. Freeman, San Francisco (1979)
27. Varasteh, A., Goudarzi, M.: Server Consolidation Techniques in Virtualized Data Centers: A Survey. IEEE Systems Journal 11, 772–783 (2017)
28. Le, K., Bianchini, R., Zhang, J., Jaluria, Y., Meng, J., Nguyen, T.D.: Reducing electricity cost through virtual machine placement in high performance computing clouds. In: 2011 International Conference for High Performance Computing, Networking, Storage and Analysis. 22:1–22:12. ACM (2011)
29. International Organization for Standardization: ISO 14044: Environmental management – Life cycle assessment – Requirements and guidelines (2006)
30. Barroso, L.A., Clidaras, J., Hölzle, U.: The Datacenter as a Computer: An Introduction to the Design of Warehouse-Scale Machines. Synthesis Lectures on Computer Architecture 8, 1–154 (2013)
31. Bouley, D.: Estimating a Data Center's Electrical Carbon Footprint. APC by Schneider Electric (2010)
32. Weber, C.: Uncertainty and variability in carbon footprinting for electronics case study of an IBM rackmount server (2011)
33. Goudarzi, H., Ghasemazar, M., Pedram, M.: SLA-based optimization of power and migration cost in cloud computing. In: Proceedings of the 2012 12th IEEE/ACM International Symposium on Cluster, Cloud and Grid Computing (ccgrid 2012). 172–179. IEEE Computer Society (2012)
34. Grinshpan, L.: Solving enterprise applications performance puzzles: queuing models to the rescue. John Wiley & Sons (2012)
35. Müller, H., Bosse, S., Turowski, K.: Optimizing server consolidation for enterprise application service providers. In: Proceedings of the 2016 Pacific Asia Conference on Information Systems (2016)
36. CO2-Bepreisung in Deutschland: Ein Überblick über Handlungsoptionen und ihre Vor- und Nachteile. Umweltbundesamt der Bundesrepublik Deutschland (2019)
37. Mi, H., Wang, H., Yin, G., Zhou, Y., Shi, D., Yuan, L.: Online Self-Reconfiguration with Performance Guarantee for Energy-Efficient Large-Scale Cloud Computing Data Centers. In: IEEE International Conference on Services Computing (SCC). 514–521. IEEE, Miami, FL, USA (2010)
38. Stillwell, M., Schanzenbach, D., Vivien, F., Casanova, H.: Resource Allocation Algorithms for Virtualized Service Hosting Platforms. Journal of Parallel and Distributed Computing 70, 962–974 (2010)
39. Wu, L., Garg, S.K., Buyya, R.: SLA-based resource allocation for software as a service provider (SaaS) in cloud computing environments. In: Proceedings of the 2011 11th IEEE/ACM International Symposium on Cluster, Cloud and Grid Computing. 195–204. IEEE Computer Society (2011)
40. Xu, J., Fortes, J.A.B.: Multi-objective Virtual Machine Placement in Virtualized Data Center Environments. In: IEEE/ACM International Conferenc on Cyber, Physical and Social Computing (CPSCom). IEEE, Hangzhou, China (2010)

https://doi.org/10.30844/wi_2020_m2-bosse