

Compensating Balancing Demand by Spatial Load Migration – The Case of Geographically Distributed Data Centers

Abstract

The increasing share of renewables confronts existing power grids with a massive challenge, stemming from additional volatility to power grids introduced by renewable energy sources. This increases the demand for balancing mechanisms, which provide balancing power to ensure that power supply always meets with demand. However, the ability to provide cost-efficient and eco-friendly balancing power can vary significantly between locations. Fridgen et al. (2017) introduce an approach based on geographically distributed data centers, aiming at the spatial migration of balancing power demand between distant locations. Although their approach enables the migration of balancing demand to cost-efficient and/or eco-friendly balancing mechanisms, it will come up against limits if deployed on a global scale. In this paper, we extend Fridgen et al. (2017)'s approach by developing a model based on geographically distributed data centers, which not only enables the migration of balancing demand but also compensates for this migration when it is contradictory between different balancing power markets without burdening conventional balancing mechanisms. Using a simulation based on real-world data, we demonstrate the possibility to exploit the potential of compensation balancing demand offered by spatial load migration resulting in economic gains that will incentivize data center operators to apply our model.

Keywords:

Distributed data centers, spatial load migration, balancing power compensation, economic potential, simulation-based case study, demand response

1. Motivation

In power grids, power generation must always meet power demand (Müller and Rammerstorfer, 2008; Rammerstorfer and Wagner, 2009). Since demand fluctuates, and is not precisely known a priori (Flinkerbusch and Heuterkes, 2010), power generation has so far been regularly adjusted in order to ensure power grid stability. However, not all energy generation types are equally appropriate to adjust their output in order to contribute to grid stability. The key challenge that accompanies generation units based on renewables (e.g. wind and solar power) is that these units are highly dependent on external factors, such as weather conditions. Especially in developed European countries, this challenge meets power systems, which are commonly designed unidirectional and top-down oriented (Palensky and Dietrich, 2011). The shift towards bidirectional electricity flows induced by renewables is associated with an increased overall complexity to the stability of the power grid that grid operators have to deal with. Current distribution and transmission grids are typically not designed to cope with the increasing feed-in of volatile renewables and the associated challenges. Consequently, the increasing integration of renewables introduces additional uncertainty to the supply side (Vandezande et al., 2010), and thus reduces the supply side's potential to balance fluctuations in demand. One possible way to counter this additional uncertainty is to use demand-side flexibility (DSF). Typical examples of the use of DSFs include the intelligent control of electric vehicle charging processes (Fridgen et al., 2014; Lujano-Rojas et al., 2012), heating and cooling systems (Ehrlich et al., 2015; Goddard et al., 2014; Grein and Pehnt, 2011), and industrial processes (Jang et al., 2016). As these examples illustrate, most of today's DSF approaches are variants on temporal flexibility (Fridgen et al., 2017). In general, the use of temporal flexibility involves rescheduling, interrupting, or omitting a power-consuming activity in order to balance the power grid.

Fridgen et al. (Fridgen et al., 2017) introduce an innovative approach which aims at spatial load migration. This approach requires a location-independent power-consuming activity which can be migrated from one location to another. One example of this kind of activity is the processing of information goods (Krcmar, 2015), which usually happens at large-scale, power-intensive data centers (DCs). In a scenario involving two geographically distributed DCs, one of the DCs could adjust its load, taking on a larger or smaller share of the DCs' overall workload and so increasing or decreasing power demand and contributing to grid stability. However, in this scenario, the other DC must then process the remaining share of the workload.

This means that the flexibility provision of the first DC also impacts the load of the second DC, resulting in unexpected load volatility. In the approach suggested by Fridgen et al. (Fridgen et al., 2017), this second DC cooperates with a local balancing mechanism – e.g., a pumped hydropower plant – which balances the unexpected load volatility. Consequently, this setup migrates the balancing demand from the balancing area of the first DC to a (possibly very distant) balancing mechanism located in the vicinity of the second DC, which delivers the required balancing power (BP). This balancing demand migration is advantageous as some markets can provide flexibility in a more cost-efficient and/or eco-friendly manner than others (Fridgen et al., 2017; Van Hulle et al., 2010; Vennemann et al., 2011).

Nevertheless, the approach suggested by Fridgen et al. (Fridgen et al., 2017) involves one major drawback: It is only concerned with the spatial migration of balancing demand to locations that can satisfy this demand in a greener or cheaper way. If their approach is applied on a global scale, the limited balancing potential of such locations would probably be overstretched.

However, the global application of the spatial BP demand migration offers another opportunity. According to Rasmussen et al. (2012), the interconnection of numerous market regions presents enormous potential, not only to migrate, but also to compensate, BP demand. This potential is based on the fact that the BP demand between different regions is largely uncorrelated, and, in some cases, even contrary, e.g., due to different weather condition (Rasmussen et al., 2012).

Furthermore, in the near future, new, omnipresent trends in information technology, such as the Internet of Things (Atzori et al., 2010) and blockchain (Schweizer et al., 2017; Swan, 2015), might further increase demand for computational power. Even more importantly to the idea of spatial compensating BP demand, these trends may further strengthen the ongoing decentralization of computational power provision. As a result, we aim to extend the work of Fridgen et al. (Fridgen et al., 2017) in order to develop an approach which virtually interconnects different BP markets, which not only allows the migration of BP demand but also its compensation without using conventional balancing mechanisms. Accordingly, this compensation allows the imbalance netting on a global scale. Thus, we describe a model that interconnects BP markets and allows them to use their contrary balancing demand; i.e., a BP market requiring positive BP is virtually connected to a BP market requiring negative BP. This results in the compensation of the balancing demand in both market areas and, thus, in a reduced need for conventional balancing mechanisms. In other words,

this paper provides a model for global imbalance netting without the use of inefficient transmission lines. The virtual interconnection of BP markets requires the participation of economic entities, which are mainly driven by economic rationale (Simon, 1979). This means that these economic entities are only likely to participate if the approach is economically feasible. Thus, the objective of this paper is:

To analyze the economic and balancing potential of virtually interconnecting different BP markets using spatial load migration based on geographically distributed DCs.

2. Demand-side flexibility based on geographically distributed DCs

In general, every DSF measure corresponds to one of the two prominent DSF concepts: load shedding and load shifting (Derakhshan et al., 2016; Feuerriegel and Neumann, 2014). Load shedding involves the abatement of a scheduled power-consuming activity, for example, switching off street lights during a power shortage (Papagiannis et al., 2008). Load shifting, on the other hand, involves the postponement of a power-consuming activity, for example, interrupting the charging process of an electric vehicle and to resume it at a later time (Fridgen et al., 2016).

The current scientific discourse extends the two existing DSF concepts, load shedding and load shifting, by introducing a third concept, load migration (Adnan et al., 2012; Fridgen et al., 2017; Wierman et al., 2014). This additional concept involves the spatial migration of a load from one location to another (Fridgen et al., 2017).

Typically, the spatial migration of load over large distances relies on high-voltage transmission lines. However, many power line construction projects fail. Reasons for their failure include excessive investments (Kishore and Singal, 2014), insecure return on investment (Buijs et al., 2011), protests by local citizens (Lütticke, 2017), and high transmission inefficiencies (Vennemann et al., 2011). A preferable, power-intensive use-case, which allows the spatial migration of load without placing additional stress on transmission lines, is a setup consisting of geographically distributed DCs (Wierman et al., 2014).

To give an example, all US-based DCs combined account for some 2% of the country's total electricity consumption, and their power demand is expected to increase (Shehabi et al., 2016). Moreover, large DC operators typically run distributed DCs, which are designed with excess capacity in order to ensure

continuous service availability (Keller and König, 2014; Zhou et al., 2010; Zissis and Lekkas, 2012). These geographically distributed DCs provide the opportunity to migrate load between DC locations by intelligently dispatching incoming requests (Ghatikar et al., 2012, 2010; Kong and Liu, 2014). Unsurprisingly, load migration enabled by geographically distributed DCs is an active research topic in the field of DC management.

3. Virtually interconnecting BP markets using distributed DCs

3.1 Setup

Our setup builds upon a strategy suggested by Fridgen et al. (Fridgen et al., 2017) involving two geographically distributed DCs, one of which participates in the local BP market and the other of which has access to a local balancing mechanism. We extend this setup by introducing a third DC (Figure 1).

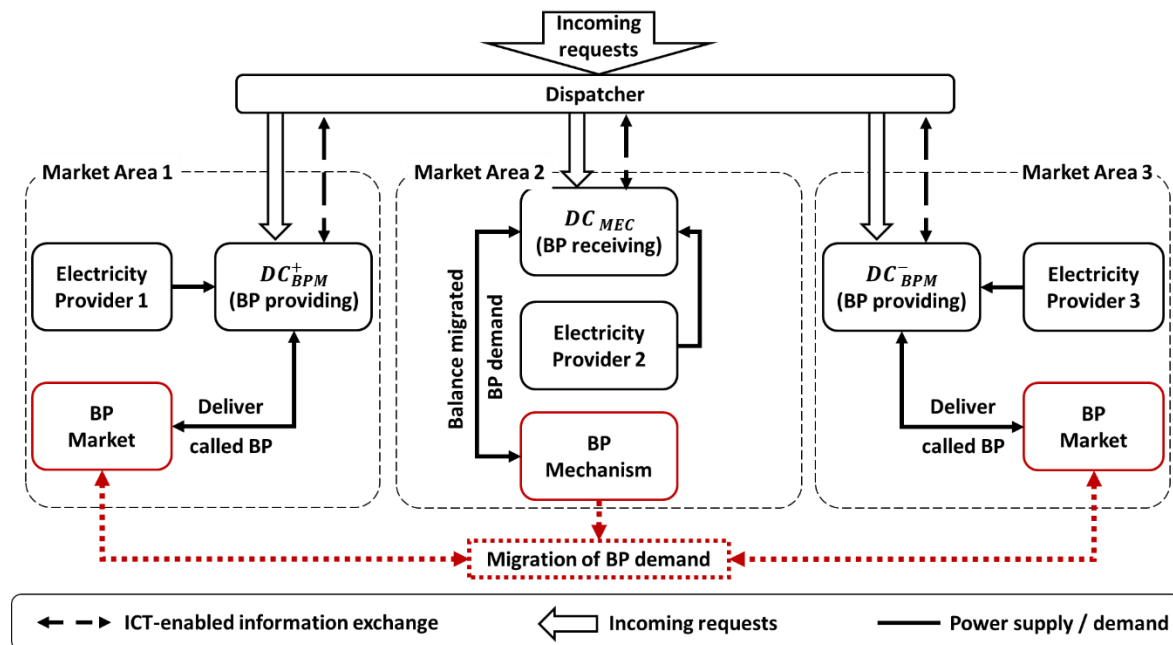


Figure 1: ICT-enabled spatial DSF approach

Each of the three DCs is located in a separate market area, e.g., in a different country or even on a different continent. In line with Fridgen et al. (Fridgen et al., 2017), all of the three DCs provide the same information processing services (Zhou et al., 2010). Accordingly, each DC is capable of processing each incoming request. In particular, large internet service providers, such as Google, Amazon, and Microsoft, operate such geographically distributed DCs (Zhang et al., 2012). The purpose of this distributed and redundant

design is to increase the availability of the provided services and to reduce the probability of data loss (Keller and König, 2014; Zissis and Lekkas, 2012). In each of the three locations, a local electricity provider supplies the DCs with power at local power prices, e.g., the power exchange prices in the respective location.

A dispatching mechanism is responsible for assigning incoming requests to the three DCs. This dispatcher is capable of regulating the DCs' utilization and, as the load of a DC largely depends on its utilization, the dispatcher can also control the load of the DCs. Two of the three DCs participate in their local BP markets. One of these two DCs participates in its local BP market in order to provide positive BP (DC_{BPM}^+). The other participates in its local BP market in order to provide negative BP (DC_{BPM}^-). Note that our model would also allow the simultaneous provision of positive and negative BP. However, for simplicity, we consider the provision of only positive or negative BP by the DCs.

In line with Fridgen et al. (Fridgen et al., 2017), we consider a market design for the two local BP markets which distinguishes between positive and negative BP. Examples of this market design include California (CAISO, 2015), Finland (Fingrid, 2017), and Germany (Consentec, 2014). Moreover, in the two local BP markets under consideration, BP can be provided by a DSF mechanism, i.e., large power consumers capable of adjusting their load. These consumers (e.g., large-scale DCs) typically report their expected load to the local electricity provider in advance (Fridgen et al., 2017).

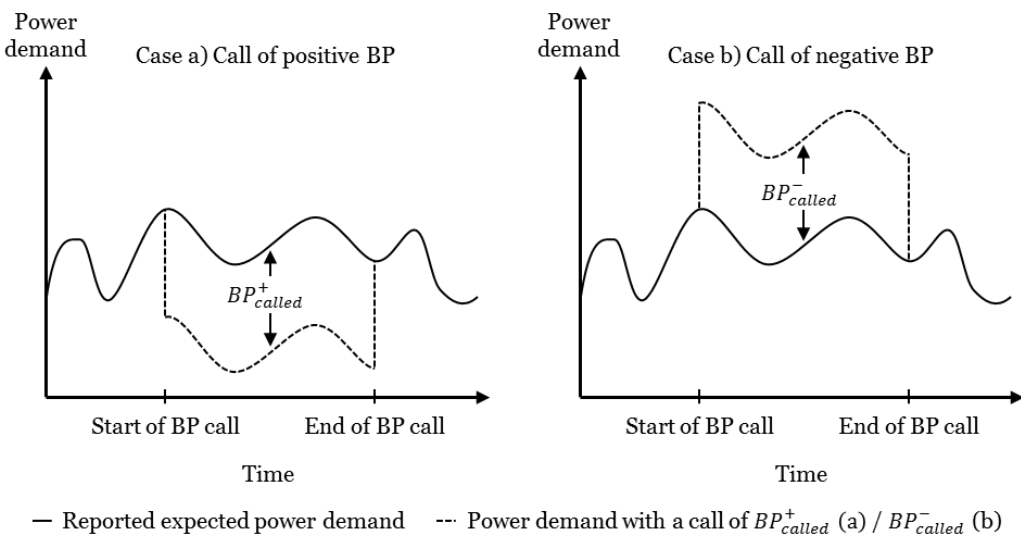


Figure 2: Providing BP via a DSF mechanism (e.g. a DC)

Whenever they are called, the consumers provide BP by deviating from their previously reported load profile. For instance, if there is a call for positive BP (i.e., if the grid is facing a deficit of electricity), the consumer can satisfy this call by decreasing their load compared to the previously reported load profile (Figure 2, case a). If there is a call for negative BP (i.e., if the grid is facing an oversupply of electricity), the consumer can satisfy this call by increasing their load compared to the previously reported load profile (Figure 2, case b). As long as the call lasts, the load must stay constant at the reported load, reduced or increased by the called amount. The delta is used as BP to stabilize the power grid.

In addition to the two DCs that can access their local BP markets, our setup involves a third DC, namely DC_{MEC} . This third DC has access to a local balancing mechanism, e.g., a pumped hydropower plant. As we illustrate in the following, DC_{MEC} is necessary to satisfy BP demand in periods when no simultaneous calls are available for compensation. In most BP markets, the participating BP providers must be able to satisfy a call for their offered BP at any point during the entire predefined period of the offer (Rebours et al., 2007a, 2007b). This makes a third DC indispensable in our setup.

3.2 Potential of DCs to offer BP

The maximum performance of a DC is defined in terms of its maximum processing capacity (MPC). The MPC describes the technically-feasible upper bound of incoming request that the DC can process within a specific period (e.g., per second). Due to their redundant design, large-scale distributed DCs typically only utilize, on average, 30% of their MPC to handle the incoming requests (Meisner et al., 2009). Therefore, DC operators switch off idle servers in order to decrease power consumption (Tolia et al., 2008). As a result, the current processing capacity (CPC), which is based on the number of running servers, is below the maximum processing capacity ($CPC \leq MPC$).

Several contributions demonstrate and further improve the proportionality between a DC's current processing capacity and its load, i.e., by switching off idle servers (Gandhi et al., 2010; Ganesh et al., 2013; Lin et al., 2013; Tolia et al., 2008; Whitney and Delforge, 2014). Drawing on these approaches, we formulate the following assumption:

Assumption 1: The load and the current processing capacity of a DC are proportional.

In this paper, we define the energy efficiency τ of a DC as the power demand required to process one request in a specific amount of time. According to our first assumption, the energy efficiency τ of a DC is independent from the DC's current processing capacity, and every incoming request generates approximately the same load.

The energy efficiency of a DC is heavily influenced by a number of different circumstances, e.g., by the installed hardware and the outdoor temperature (Avelar et al., 2012). Although these circumstances are highly individual to each DC, the information that has been published about Google's globally distributed DCs illustrates that the energy efficiency of modern DCs' varies only slightly (Alphabet Inc., 2017). Based on this information, and in the interest of simplicity, we make Assumption 2, in line with Fridgen et al. (Fridgen et al., 2017):

Assumption 2: The DCs under consideration have the same energy efficiency
 $(\tau_{DC_{BPM+}} = \tau_{DC_{BPM-}} = \tau_{DC_{MEC}})$.

This second assumption allows us to abstract from DC-specific circumstances, and thus to demonstrate the economic potential of load migration in a more general manner.

In most BP markets, the participating BP providers must be able to satisfy a call for the BP they offer at any point during the entire predefined offer period (Rebours et al., 2007a, 2007b). Consequently, the load of the DCs must always enable them to deliver the offered BP. The load of DCs, however, depends on their current processing capacity and, thus, on the number of incoming requests that the DC has to process. In general, the exact number of incoming requests is volatile and unknown a priori. Nevertheless, there are a number of scientific contributions which suggest approaches to forecast the number of incoming requests (Gmach et al., 2007; Jheng et al., 2014; Tseng et al., 2018). Drawing on these existing forecasting approaches, we formulate our third assumption:

Assumption 3: The amount of BP available for offering is known.

In addition to defining the amount of BP available for offering, the dispatcher also has to define bid prices. The remuneration methods of existing BP markets vary greatly (Rebours et al., 2007b). Consequently, several articles describe various pricing strategies for the different remuneration methods (e.g., Bajpai and Singh, 2007; Krishna, 2010; Swider, 2006; Wang et al., 2017). Since we aim to demonstrate the economic

and balancing potential of our approach, we focus on determining an upper bound to the potential by applying a pricing strategy, which is based on perfect information.

3.4 Migrating BP with geographically distributed DCs

Inspired by the approach of Fridgen et al. (Fridgen et al., 2017), in this section we describe a model which can be used to interconnect distant BP markets. For this, we suppose that the dispatcher placed in the location of DC_{BPM}^+ an offer of positive BP in the amount of BP_{bid}^+ . Consequently, during the entire offer period, the local BP market can call DC_{BPM}^+ to deliver positive BP of, at most, BP_{bid}^+ . In order to be prepared for the BP delivery, the dispatcher must always be able to migrate load in the amount of BP_{bid}^+ from DC_{BPM}^+ to DC_{MEC} by shifting the corresponding workload. Moreover, we suppose that the dispatcher placed in the location of DC_{BPM}^- an offer for negative BP in the amount of BP_{bid}^- , with $BP_{bid}^- = BP_{bid}^+$. To guarantee the delivery of the offered negative BP, the dispatcher must always be able to migrate load in the amount of BP_{bid}^- from DC_{BPM}^- to DC_{MEC} .

In the following, we illustrate our model by separately describing three different cases. In the first case, DC_{BPM}^+ is called to deliver positive BP (cf. section 3.4.1). In the second case, DC_{BPM}^- is called to deliver negative BP (cf. section 3.4.2). In the third case, there is a positive BP call for DC_{BPM}^+ and, simultaneously, a negative BP call for DC_{BPM}^- (cf. section 3.4.3).

3.4.1 Positive BP call

Provided that there is a call for positive BP in the amount of $BP_{called}^+ \leq BP_{bid}^+$, the dispatcher shifts workload equivalent to BP_{called}^+ from DC_{BPM}^+ to DC_{MEC} in order to reduce the load of DC_{BPM}^+ by the called amount of positive BP. This workload shift requires an adjustment procedure, which is summarized in Figure 3. The procedure consists of three steps and is necessary for the migration of positive BP demand (Fridgen et al., 2017):

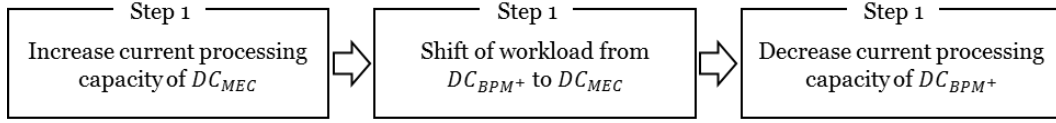


Figure 3: Positive BP demand

Step 1 of the adjustment procedure:

First of all, the dispatcher has to increase the current processing capacity of DC_{MEC} , e.g., by switching on additional servers, in order to allow the shift of workload from DC_{BPM}^+ to DC_{MEC} . Since we assume both DCs to have the same energy efficiency, the increased current processing capacity of DC_{MEC} results in an unplanned and, thus, unreported increase to the load in terms of the amount of BP_{called}^+ . This unplanned additional load results in an additional demand for BP. To ensure that this does not destabilize the power grid in the location of DC_{MEC} , the dispatcher simultaneously calls the local balancing mechanism to cover the additional BP demand.

Step 2 of the adjustment procedure:

As soon as the additional processing capacity of DC_{MEC} is available (e.g., as soon as the additional servers are fully functional), DC_{MEC} is able to process more incoming requests. The dispatcher uses this additional capacity by assigning fewer requests to DC_{BPM}^+ and more to DC_{MEC} . As a result, the workload of DC_{BPM}^+ decreases and the workload of DC_{MEC} increases.

Step 3 of the adjustment procedure:

After shifting the workload from DC_{BPM}^+ to DC_{MEC} , the dispatcher reduces the current processing capacity of DC_{BPM}^+ (e.g., by switching off idle servers). By reducing the current processing capacity, the dispatcher also reduces the load of DC_{BPM}^+ , resulting in a negative deviation from the DC's planned and reported load, i.e., the demand is lower than reported. This difference between the reported and actual load is passed on to the power grid as the called positive BP.

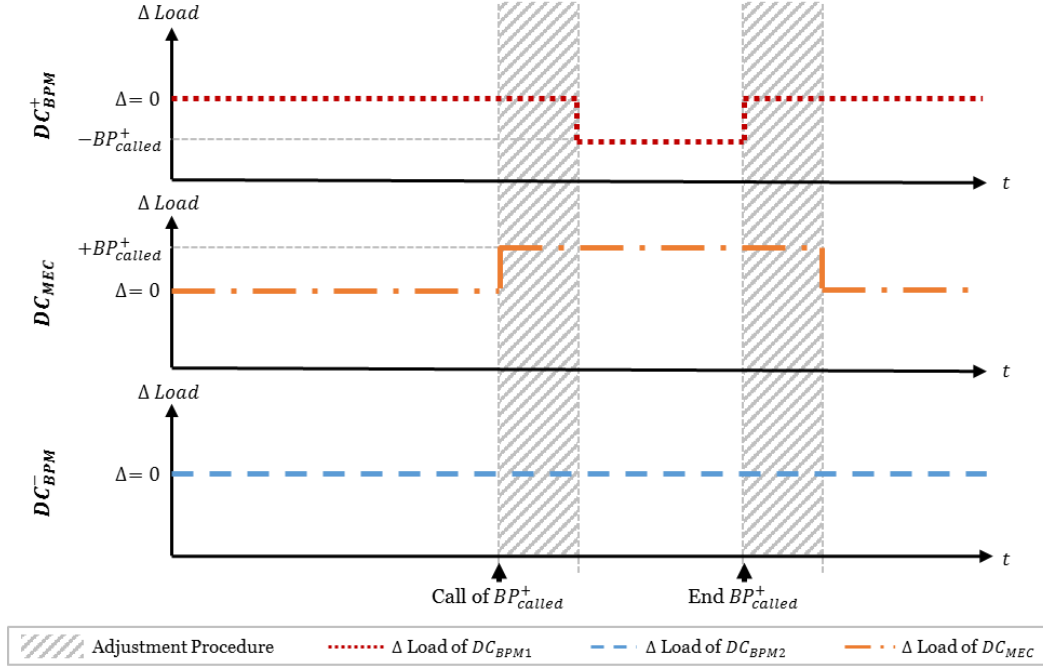


Figure 4: Adjustment procedure for positive BP

Consequently, whenever there is a call for positive BP, the load of DC_{MEC} immediately increases at the beginning of the adjustment procedure, while the load of DC_{BPM}^+ decreases only at the end. With the end of the third step, the adjustment procedure terminates. Until the call for positive BP ends again, DC_{BPM}^+ delivers the called BP to the local BP market, and DC_{MEC} receives the corresponding positive BP from a local BP mechanism, which covers the migrated BP demand (Figure 4). When the call ends, the dispatcher starts a new adjustment procedure by performing in reverse the three steps described above; i.e., firstly, it increases the current processing capacity of DC_{BPM}^+ , secondly, it shifts the corresponding workload from DC_{MEC} to DC_{BPM}^+ , and, thirdly, it decreases the current processing capacity of DC_{MEC} .

3.4.2 Negative BP call

In contrast to the previous section, we suppose here that DC_{BPM}^- has to satisfy a call for negative BP in the amount of $BP_{called}^- \leq BP_{bid}^-$. Accordingly, the dispatcher must increase the load of DC_{BPM}^- by BP_{called}^- in order to counteract this electricity surplus. To do so, the dispatcher has to perform in reverse the process of migrating positive BP. First of all, the dispatcher increases the current processing capacity of DC_{BPM}^- equivalent to BP_{called}^- in order to enable the shift of the workload (step 1 of the adjustment procedure). Due

to the increasing current processing capacity, the load of DC_{BPM}^- increases right at the beginning of the adjustment procedure and, thus, immediately provides the called negative BP (Figure 5).




Figure 5: Adjustment procedure for negative BP

The second step of the adjustment procedure is to shift workload from DC_{MEC} to DC_{BPM}^- in order to use the additional processing capacity of DC_{BPM}^- . In the third step, the dispatcher reduces the current processing capacity of DC_{MEC} (e.g., by switching off idle servers). Since we assume the energy efficiency of both DCs to be equal, this third step reduces the load of DC_{MEC} by BP_{called}^- .

Consequently, whenever there is a call for negative BP, the load of DC_{BPM}^- immediately increases at the beginning of the adjustment procedure, and the load of DC_{MEC} decreases at the end. With the end of the third step, the adjustment procedure terminates. Until the call ends again, DC_{BPM}^- delivers the called BP to the local BP market and DC_{MEC} receives the corresponding negative BP from a local BP mechanism, which covers the migrated BP demand. When the call ends, the dispatcher starts a new adjustment procedure by performing in reverse the steps described above.

3.4.3 Simultaneous positive and negative BP call

In this section, we illustrate a case in which a positive BP call in the location of DC_{BPM}^+ and a negative BP call in the location of DC_{BPM}^- occur simultaneously. This third case, then, is the combination of the first two, namely a positive BP call and a negative BP call. For clarity, we suppose in this section that the called amount of positive BP in the location of DC_{BPM}^+ is equal to the called amount of negative BP in the location of DC_{BPM}^- , i.e. $BP_{called}^+ = BP_{called}^-$. Moreover, we suppose that the positive call starts and ends during the negative call. Note that the model also applies if the called amount differs in both locations, and if the calls start and/or end in a different order.

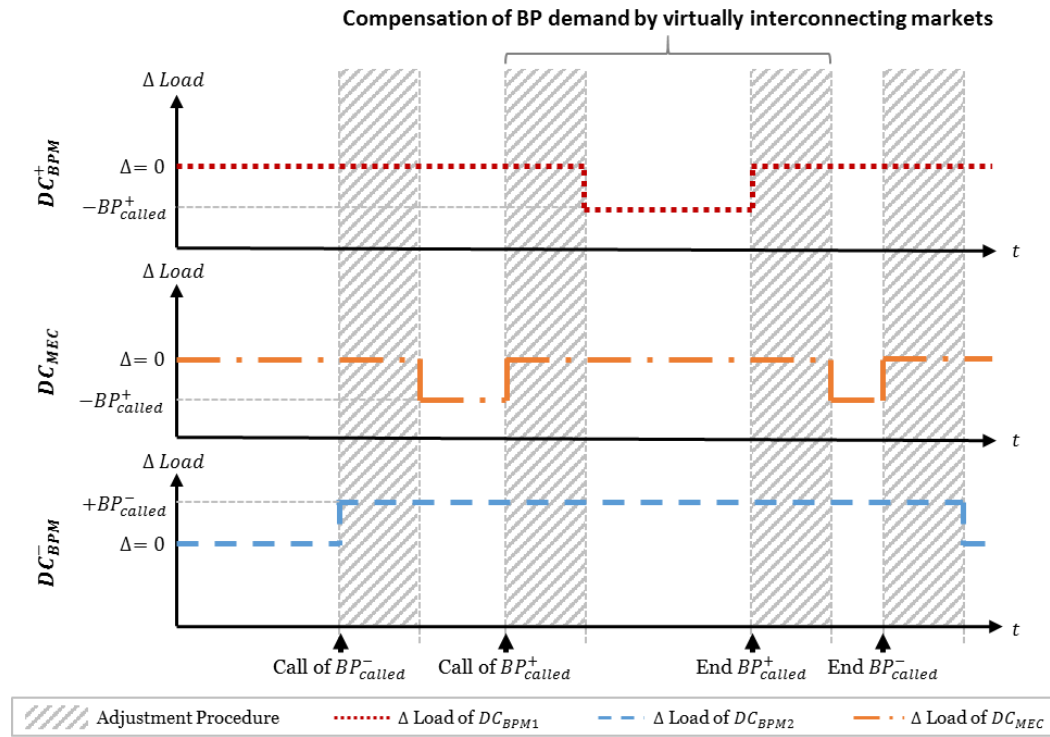


Figure 6: Adjustment procedure for contrary BP demand

Figure 6 illustrates the two calls described above. As soon as the negative BP is called, the dispatcher initiates the adjustment procedure by increasing the current processing capacity of DC_{BPM}^- . As described in the previous section, this immediately results in an increasing load of DC_{BPM}^- by BP_{called}^- and, thus, the required negative BP is delivered. At the end of the adjustment procedure, the load of DC_{MEC} is reduced by BP_{called}^- . Consequently, when the adjustment procedure terminates, the load of DC_{BPM}^- is reduced and the load of DC_{MEC} is increased by BP_{called}^- .

As soon as the BP market calls DC_{BPM}^+ to deliver positive BP, the dispatcher initiates a second adjustment procedure in order to migrate the positive BP demand from the location of DC_{BPM}^+ to the location of DC_{MEC} . Right at the beginning of this adjustment procedure, the dispatcher increases the current processing capacity of DC_{MEC} to allow the shift of workload from DC_{BPM}^+ to DC_{MEC} . This results in an additional load of DC_{MEC} by the amount of BP_{called}^+ , which is held until the call ends and the current processing capacity is reduced again. However, in response to the negative BP call, the dispatcher has already reduced the load of DC_{MEC} by BP_{called}^- . Thus, a simultaneous positive and negative call results in compensation of the BP demand in the location DC_{MEC} . Since we assume in this example $BP_{called}^+ = BP_{called}^-$, the BP demand is fully compensated and, thus, neither positive nor negative BP delivered by a conventional balancing mechanism is necessary to satisfy the two BP calls.

Consequently, whenever there are contrary BP calls in the two BP markets, the dispatcher compensates by intelligently migrating the BP demand. Accordingly, the two BP markets are virtually interconnected via a third DC location, and the contrary BP demand is used to simultaneously stabilize the two markets without the use of a conventional balancing mechanism. If several DCs all over the world, rather than only three, participate in this BP-market interconnection model, the balancing potential could be massive due to uncorrelated BP demand around the world (Chatzivasileiadis et al., 2013; Gellings, 2015).

4. Evaluation of the BP migration's economic and balancing potential

To analyze the economic and balancing potential of virtually interconnecting distant BP markets, we decided to conduct a comprehensive simulation study based on real-world data.

4.1 Simulation setting

We instantiate the above-described setup with the market regions of three different countries in Europe, namely Germany, Finland, and the Netherlands. In all three countries there are separate markets for trading BP, whereby these markets use auction mechanisms based on the merit-order principle to decide which BP supplier is called to deliver BP. Moreover, the BP markets in the three countries differ between positive and negative BP. Each BP market allows to trade predefined standardized BP products which determine parameters as the reaction time (i.e., the allowed period between the initial call signal and the

actual BP delivery) and the minimum and maximum duration of a call. Accordingly, these three countries are used as examples in our evaluation because of their suitable market structure, comparable BP products, and data availability. However, other countries, even with a completely different market structure, could also participate.

For this evaluation, we suppose that a central dispatcher assigns incoming requests to three geographically distributed DCs; one in Germany, one in Finland, and one in the Netherlands. The German and Dutch DCs participate in the respective local BP markets, while the Finish CS is connected to a balancing mechanism. We take into consideration the virtual interconnection of the German and the Dutch BP market. In our evaluation we provide four different evaluation scenarios based on the combination of two different pricing strategies and two different volume strategies that a DC operator can chose from. The combination of the two different volume strategies and the two different pricing strategies results in four evaluation scenarios (Figure 7). To provide a reference case, we also consider a fifth strategy that does not consider BP.

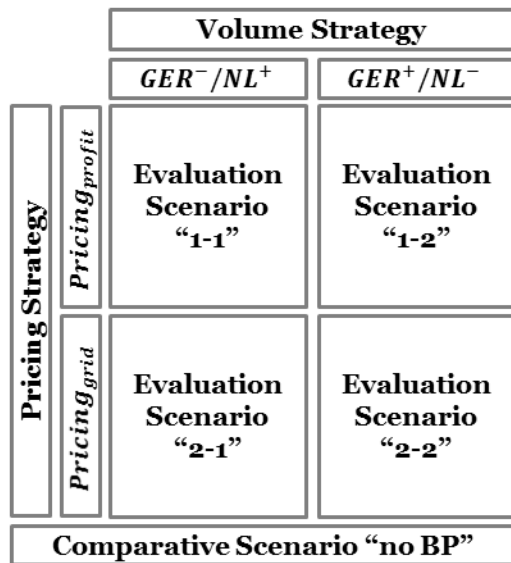


Figure 7: Evaluation scenarios

The volume strategies available to the DC operator determine the type of BP offered in the German and Dutch BP markets. Since the amount of BP that DCs can offer varies greatly because of numerous influencing factors (e.g., the size of the DC), we standardize the bidden amount of BP to 1 MW, which the DCs can deliver during the entire offer period (Fridgen et al., 2017). Based on this standardization, we define two volume strategies:

- GER^-/NL^+ : The dispatcher offers 1 MW of negative BP in Germany and 1 MW of positive BP in the Netherlands
- GER^+/NL^- : The dispatcher offers 1 MW of positive BP in Germany and 1 MW of negative BP in the Netherlands

The DC operator chooses a pricing strategy. The pricing strategy, in turn, defines the bid prices in both locations. The selected pricing strategy is important in both locations, since in Germany and in the Netherlands the BP markets call the BP providers based on a merit order regarding the individually bidden prices. Both applied pricing strategies ($Pricing_{profit}$ and $Pricing_{grid}$) are based on perfect information about the past, the present, and the future. Thus, our dispatcher can bid optimal prices.

- $Pricing_{profit}$: The dispatcher offers BP at optimal prices, which maximize the monetary value and, thus, the profit to be made by virtually interconnecting the two BP markets of Germany and the Netherlands
- $Pricing_{grid}$: The dispatcher offers BP at the optimal prices, which maximizes the number of simultaneous calls in Germany and the Netherlands and, thus, maximizes the potential to compensate BP demand

In a fifth evaluation scenario, the dispatcher offers no BP and, thus, assigns the incoming requests to the DCs, which, can process them, on average, at the lowest cost. Consequently, this scenario represents a comparative scenario, which is necessary to determine the economic and balancing potential of the other four scenarios. We assess the economic feasibility by determining the cost and benefit difference between each of the other four evaluation scenarios compared to the fifth, comparative scenario “no BP”. We include in this analysis the costs of the power required to process the workload, which is necessary to migrate the offered BP. Additionally, we include the benefits of offering and delivering BP in Germany and the Netherlands. We also consider the opportunity costs for the balancing mechanism in Finland. We define these opportunity costs in terms of the potential remuneration a balancing mechanism would receive when delivering the called BP in Finland rather than migrating it to one of the other two countries (Fridgen et al., 2017).

4.2 Simulation data set

For this paper's evaluation, we developed a time-discrete simulation model. To define this time-discrete simulation model, we discretized all input parameters on a five-minute basis. A finer discretization would only result in apparent accuracy due to data availability. A coarser discretization, on the other hand, would lead to inaccuracies as the considered type of BP is called in the range of about five to 15 minutes. We implemented the simulation model in the object-oriented programming language Java. As the basis of our simulation environment, we embedded MASON which is an open-source library for discrete-event multi-agent simulations.

With the time-discrete simulation model, we simulated the weekly economic value of our model for the entire year of 2017 in order to compensate for seasonal effects. To arrive at a realistic estimation of the economic potential, we considered the real-world BP offers, BP demand, and BP prices in 2017 for the secondary BP market in Germany and the regulating power market in the Netherlands. These two markets have similar structures and, thus, are well-suited to virtual interconnection. The data for the German BP market is published by the union of the four German transmission system operators (regelleistung.net, 2018). The data for the Dutch BP markets is published by TenneT (TenneT, 2018). In line with the data resolution of both BP markets, we include in our simulation the BP demand with a granularity of 15 minutes. Moreover, in both BP markets, the response time between call and the actual full BP delivery is 5 minutes (Consentec, 2014; TenneT, 2012). However, for both markets, there is no information available on the actual duration of the individual calls which the BP providers have to satisfy. Thus, we follow Fridgen et al. (Fridgen et al., 2017) in that assuming the calls for BP are equally distributed between 5 and 15 minutes. To define opportunity costs for the balancing mechanism, we once again follow Fridgen et al. (Fridgen et al., 2017) in using the Finnish BP prices derived from Nord Pool Spot (2018). For the local power prices in all three areas, we use hourly spot market prices from the respective power exchanges (ENTSO-E, 2018; Nord Pool Spot, 2018).

For the three DCs, we consider the adjustment phase required to startup and shutdown servers in order to provide the called BP. Gandhi et al. (2010) suggests that the startup time of servers is 200 seconds and the shutdown delay is negligible. Due to ongoing hardware and software innovations, the adjustment time for starting-up servers has steadily decreased in the recent years and is now less than one minute (Mao and

Humphrey, 2012; Qu et al., 2017). We also follow Gandhi et al. (2010) in considering the processing time of incoming requests to be negligibly short. For the purpose of our simulation, we set the duration of the adjustment phase to a discretionary period of one minute which matches to our simulation incremental time step. All input variables from our simulation are summarized in Table 1.

Table 1: Simulation input

<i>Description</i>	<i>Value</i>	<i>Source</i>
German BP offers, calls, and prices	Real-world data	Regelleistung.net (2018)
Dutch BP offers, calls, and prices	Real-world data	TenneT (2018)
Duration of BP calls	Assumed Uniform Distribution (5 min, 15 min)	Consentec (2014)
Finnish BP prices	Real-world data	Nord Pool Spot (2018)
German power prices	Real-world data	ENTSO-E (2018)
Dutch power prices	Real-world data	ENTSO-E (2018)
Finnish power prices	Real-world data	Nord Pool Spot (2018)
Duration of the adjustment phase	$d_{adjust} = 1 \text{ min}$	Gandhi et al. (2010), Mao and Humphrey (2012), Qu et al. (2017)

4.3 Simulation results

Figure 8 compares the yearly economic and compensation potential of all four evaluation scenarios, whereby the former two scenarios are based on the pricing strategy $Pricing_{profit}$ and the latter two scenarios are based on the pricing strategy $Pricing_{grid}$.

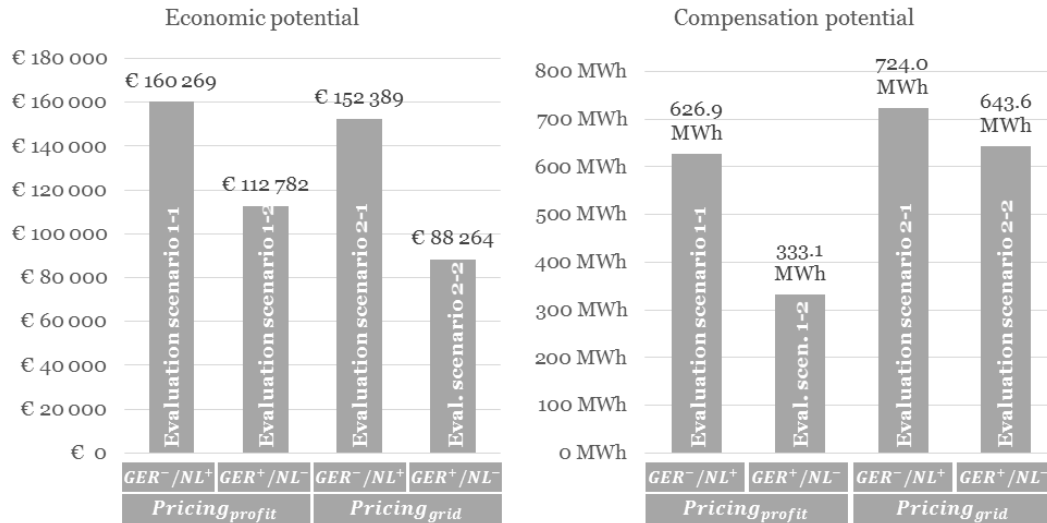


Figure 8: Yearly economic and BP compensation potential (1 MW of BP in both BP market locations)

The results of the four, year-long evaluation scenarios point to three different key findings:

- Economic potential** (Figure 8 left): In all four evaluation scenarios, the interconnection of distant BP markets results in positive economic value. Thus, in all evaluation scenarios, the DC setup is incentivized to participate in the two balancing markets and to provide BP by shifting workload. However, the economic value achievable using the pricing strategy *Pricing_{profit}* (€ 160 269 and € 112 782) is higher than that achievable with the pricing strategy *Pricing_{grid}* (€ 152 389 and € 88 264). This is reasonable, since the second pricing strategy results in BP bids which maximize the compensation potential, even if this results in lower economic potential. Regardless of the pricing strategy employed, the volume strategy *GER⁻/NL⁺* results in a higher economic value (€ 160 269 and € 152 389) than the volume strategy *GER⁺/NL⁻* (€ 112 782 and € 88 264). Accordingly, due to the results in Figure 8, it is, on average, economically favorable to provide 1 MW of negative BP in Germany and 1 MW of positive BP in the Netherlands.
- Compensation of BP demand** (Figure 8 right): In all four evaluation scenarios, the interconnection of BP markets results in the compensation of BP demand. Independent of the selected volume and pricing strategy, the setup contributes to grid stability in both Germany and the Netherlands. However, the amount of compensated BP demand differs between the four evaluation scenarios. In evaluation scenario 1-1 (evaluation scenario 1-2) the compensated BP demand is lower than in evaluation scenario 2-1 (evaluation scenario 2-2). Accordingly, when using

pricing strategy $Pricing_{profit}$, the compensated BP demand is lower than when using the pricing strategy $Pricing_{grid}$. This is reasonable, since the pricing strategy $Pricing_{grid}$ aims to maximize the compensated BP demand. However, the discrepancy in the BP compensation potential illustrates that the current level of remuneration for BP demand does not fully incentivize the compensation of BP. Moreover, volume strategy GER^-/NL^+ offers a higher compensation potential than volume strategy GER^+/NL^- . As a consequence, due to Figure 8, the highest compensation can be achieved by using pricing strategy $Pricing_{grid}$ and volume strategy GER^-/NL^+ which results in a potential of 724.0 MWh.

- **Dominance of volume strategy GER^-/NL^+ :** Providing that the DC setup cannot switch between the two volume strategies during the simulated year, the volume strategy GER^-/NL^+ dominates the volume strategy GER^+/NL^- . This means that, regardless of the applied pricing strategy, the economic and compensation potential of the GER^-/NL^+ volume strategy exceeds the potential of the GER^+/NL^- strategy. Accordingly, in our evaluation, the defined DC setup should provide 1 MW of negative BP in Germany and 1 MW of positive BP in the Netherlands. Note that intelligently switching between both volume strategies during the year would presumably further increase the economic and BP compensation potential. This is because the potential of the two volume strategies varies during the evaluated year, as indicated in the evaluation scenarios 1-1 and 1-2 by figures 9 and 10. However, the combination of the two volume strategies would increase the complexity for the DC operator. Therefore, we exclude in our evaluation the option to switch and leave this as a subject for further research.

Besides the volatility of the economic and compensation potential for the two evaluation scenarios 1-1 and 1-2, Figures 9 and 10 illustrate that the economic potential in the individual periods is not necessarily accompanied with a high compensation potential. In evaluation scenario 1-2 (Figure 10), for instance, there is a low economic potential (red dashed line) and a high compensation potential (grey bars) achievable in period 7 and in period 43 vice versa. However, there are also periods, in which a high economic potential is combined with a high compensation potential (e.g. period 4 in evaluation scenario 1-1 illustrated in Figure 9 and period 12 in evaluation scenario 1-2 illustrated in Figure 10). This is due to the fact that the economic and the compensation potential largely depends on the remuneration for providing BP. Thus, the offers of

other BP providers in the two locations have a high impact and can significantly vary between the offer periods. Accordingly, the maximization of the economic and/or compensation potential can profit from the in-depth analysis of the competitive situation at the BP markets.

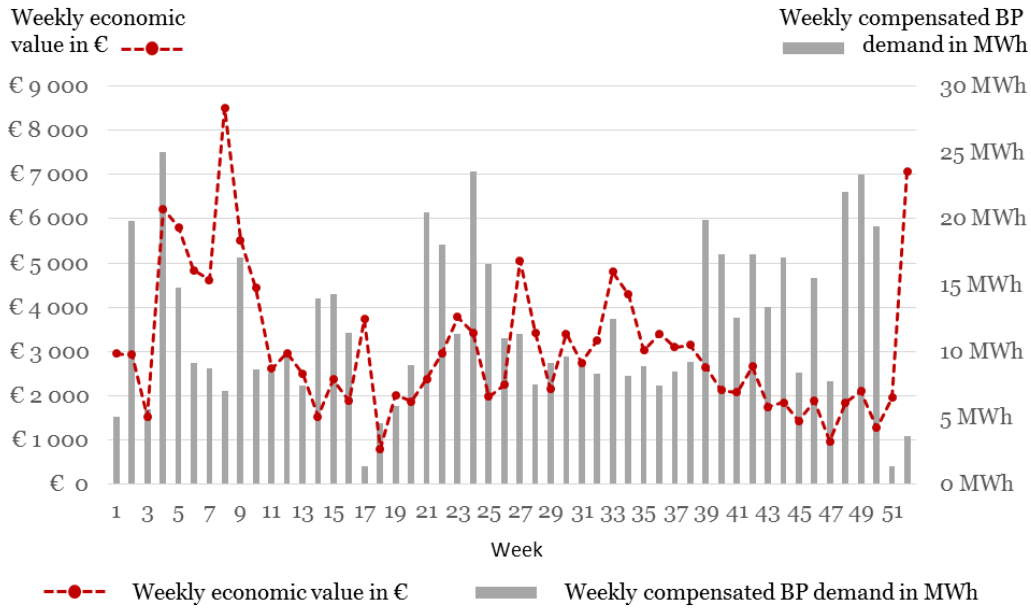


Figure 9: Weekly economic value and weekly compensated BP demand of evaluation scenarios 1-1

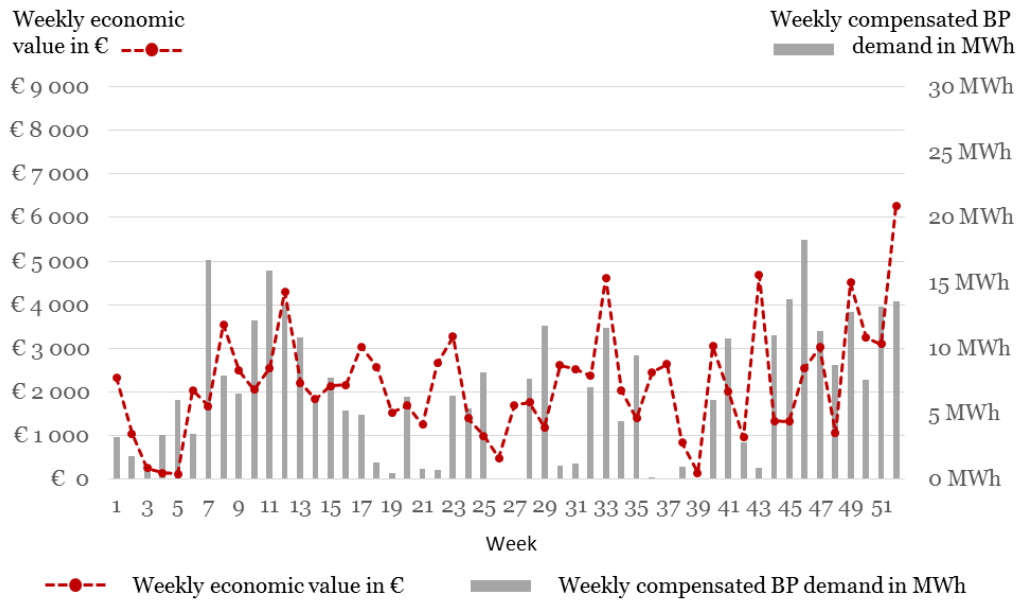


Figure 10: Weekly economic value and weekly compensated BP demand of evaluation scenarios 1-2

To test the results for robustness, we conducted 50 simulation runs for each evaluation scenario. For each scenario, Table 2 summarizes the minimum, the lower quartile, the median, the upper quartile, and the maximum values resulting from these 50 simulation runs.

Table 2: Yearly results retrieved from 50 simulation runs (1 MW of BP in both BP market locations)

	<i>Pricing Strategy</i>	<i>Volume Strategy</i>	<i>Minimum</i>	<i>Lower Quartile</i>	<i>Median</i>	<i>Upper Quartile</i>	<i>Maximum</i>
<i>Economic potential</i>	<i>Pricing_{profit}</i>	<i>GER⁻/NL⁺</i>	€ 159 403	€ 159 932	€ 160 269	€ 160 699	€ 161 561
		<i>GER⁺/NL⁻</i>	€ 112 130	€ 112 552	€ 112 782	€ 112 999	€ 113 317
	<i>Pricing_{grid}</i>	<i>GER⁻/NL⁺</i>	€ 150 913	€ 152 047	€ 152 389	€ 152 675	€ 153 322
		<i>GER⁺/NL⁻</i>	€ 87 459	€ 87 979	€ 88 264	€ 88 535	€ 88 852
<i>Compensation potential</i>	<i>Pricing_{profit}</i>	<i>GER⁻/NL⁺</i>	614 MWh	624 MWh	627 MWh	634 MWh	645 MWh
		<i>GER⁺/NL⁻</i>	332 MWh	330 MWh	333 MWh	336 MWh	349 MWh
	<i>Pricing_{grid}</i>	<i>GER⁻/NL⁺</i>	712 MWh	721 MWh	724 MWh	729 MWh	735 MWh
		<i>GER⁺/NL⁻</i>	633 MWh	641 MWh	644 MWh	647 MWh	651 MWh

Within the 50 simulation runs, the values for the economic potential fluctuate by less than one percent around the median. For the compensation potential, the values fluctuate by less than five percent around the median. This low variability of the results confirms that the stochastic simulation input is largely compensated due to the long period of simulation (i.e. the simulation of an entire year). Thus, the results of Table 2 suggest that conducting more than 50 simulation runs would not result in fundamentally new or other key findings.

Summarizing the above, our evaluation illustrates that there is economic incentive for a setup consisting of three DCs in three different market regions, which can provide its spatial flexibility as BP. In our evaluation, we standardized the amount of BP available for bidding at 1 MW. To give an estimation of the economic and the BP compensation potential for a real-world setup, we ran our simulation model with input data relevant to a real-world Google DC setup which is inspired by Fridgen et al. (Fridgen et al., 2017). Google operates several large, cloud-scale DCs worldwide. Three of Google's DCs are located in Hamina (Finland), Eemshaven (Netherlands), and St. Ghislain (Belgium). Since Google does not operate a DC in Germany, but as the Belgian DC is only about 150 kilometers away from the German border, we count this DC as being located in Germany.

Google does not publish information about the energy consumption of the three DCs. According to Hintermann und Clausen (2014), the power consumption of large-scale DCs is in the region of 50 MW. Koomey and Taylor (2015) point out that, even in modern, large-scale DCs, approximately 30% of the maximum processing capacity is comatose, meaning that the corresponding servers "have delivered no information or computation services in 6 months or more". In line with Fridgen et al. (Fridgen et al., 2017), we thus suppose that one of Google's large-scale DCs is capable of providing at least 15 MW of BP via spatially load migration. Accordingly, in volume strategy GER^-/NL^+ , the DC setup simultaneously offers 15 MW of negative BP in Germany and 15 MW of positive BP in the Netherlands. In volume strategy GER^+/NL^- , however, the DC setups simultaneously offers 15 MW of positive BP in Germany and 15 MW of negative BP in the Netherlands.

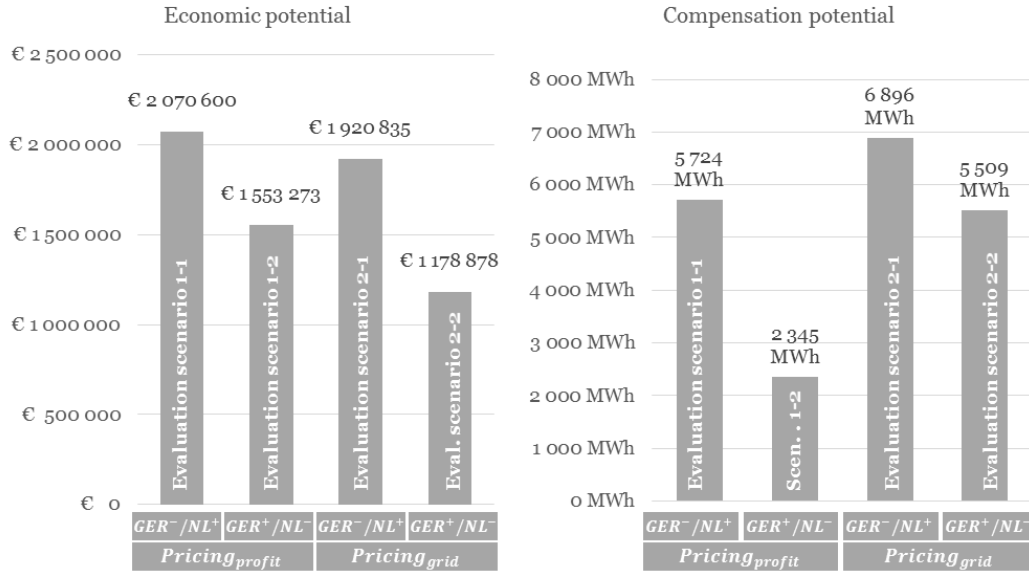


Figure 11: Yearly economic and BP compensation potential (15 MW of BP in both BP market locations)

The increase in the level of BP provided (from 1 MW to 15 MW in each of the two locations) results in an almost proportionally-increased, yearly economic value for all four evaluation scenarios (Figure 11). For the Google DC setup, the volume strategy GER^-/NL^+ dominates the volume strategy GER^+/NL^- , meaning that the economic and compensation potential is greater for the former volume strategy, irrespective of the pricing strategy used. Based on our evaluation of the featured Google DC setup, applying the pricing strategy $Pricing_{profit}$ could reach a yearly economic potential of over € 2 million, and compensate more than 5 GWh of BP in both locations by virtually interconnecting the German and the Dutch BP markets. The use of the second pricing strategy, $Pricing_{grid}$, results in an even higher BP demand compensation potential of almost 7 GWh, which is incentivized by a smaller but substantial economic potential of over € 1.9 million. Again, for all four evaluation scenarios, the results of the 50 simulation runs illustrate that variability is low and, thus, conducting more than 50 runs seems unnecessary (Table 3).

Table 3: Yearly results retrieved from 50 simulation runs (15 MW of BP in both BP market locations)

	<i>Pricing Strategy</i>	<i>Volume Strategy</i>	<i>Minimum</i>	<i>Lower Quartile</i>	<i>Median</i>	<i>Upper Quartile</i>	<i>Maximum</i>
<i>Economic potential</i>	<i>Pricing_{profit}</i>	<i>GER⁻/NL⁺</i>	€ 2 048 634	€ 2 065 433	€ 2 070 600	€ 2 073 307	€ 2 089 596
		<i>GER⁺/NL⁻</i>	€ 1 544 522	€ 1 550 301	€ 1 553 273	€ 1 556 171	€ 1 562 682
	<i>Pricing_{grid}</i>	<i>GER⁻/NL⁺</i>	€ 1 911 268	€ 1 916 479	€ 1 920 835	€ 1 925 560	€ 1 939 827
		<i>GER⁺/NL⁻</i>	€ 1 171 396	€ 1 176 340	€ 1 178 878	€ 1 181 845	€ 1 187 778
<i>Compensation potential</i>	<i>Pricing_{profit}</i>	<i>GER⁻/NL⁺</i>	5 541 MWh	5 666 MWh	5 724 MWh	5 781 MWh	5 865 MWh
		<i>GER⁺/NL⁻</i>	2 248 MWh	2 305 MWh	2 345 MWh	2 371 MWh	2 442 MWh
	<i>Pricing_{grid}</i>	<i>GER⁻/NL⁺</i>	6 780 MWh	6 865 MWh	6 896 MWh	6 939 MWh	7 035 MWh
		<i>GER⁺/NL⁻</i>	5 345 MWh	5 476 MWh	5 509 MWh	5 565 MWh	5 618 MWh

To summarize, there is an economic incentive for both the standardized setup and the Google DC setup to provide their DSF potential by virtually interconnecting BP markets. By doing so, the DCs allow for the compensation of BP demand without using a conventional balancing mechanism.

5. Conclusion and policy implications

5.1 Summary and Policy implications

Our paper illustrates the economic and balancing potential of *virtually interconnecting different BP markets by spatial load migration based on geographically distributed DCs*. However, we do not only give an initial estimation of the economic potential, but also describe the resulting policy-related implications.

DSF can be principally differentiated between temporally or spatially utilization of flexibility Fridgen et al. (2017). This differentiation is also possible for the flexibility of DCs (Keller et al., 2019). Temporal flexibility might significantly contribute to the stability of the power grid of the future (Schott et al., 2018). However, approaches on temporal flexibility use the rescheduling, interrupting, or omitting of a power-consuming activity to provide the flexibility (Fridgen et al., 2017). Whether it is in the field of research of flexible demand in the energy sector in general (e.g., Ehrlich et al., 2015; Fridgen et al., 2015, 2014; Goddard et al.,

2014; Grein and Pehnt, 2011; Jang et al., 2016; Lujano-Rojas et al., 2012) or in the field of DCs in special (e.g., Andrzejak et al., 2010; Keller et al., 2019; Vieira et al., 2015) both use the the rescheduling, interrupting or omitting for the provision of flexibility in energy demand (Fridgen et al., 2017). In the field of rescheduling or interrupting in the domain of DCs only few requests are flexible as well as delay-tolerant (Gmach et al., 2010). The possibility to provide DSF by omitting typically comes with quality degradation under the constraints from quality-of-service requirements and service level agreements as outlined by (Wierman et al., 2014). These changes in the temporal overall demand can be avoided by spatially shifting the power consuming activity to another DC (Fridgen et al., 2017; Müller and Rammerstorfer, 2008; Rebours et al., 2007). Our contribution therefore extends the previously introduced concept of Fridgen et al. (2017). Their model allows for interconnecting different BP markets via the geographic shifting of a power consuming activity (i.e., the processing of data) between different DC location using a balancing mechanism. However, there is a potential major drawback of spatial load migration, as collaboration of balancing mechanisms, e.g. pumped hydropower plants is required to avoid that one power grid improves its stability at the expense of another. Accordingly, our contribution extends the basic idea in this paper by connecting at least two physical market places for BP. When BP is provided only in one market area a balancing mechanism to compensate for the spatial migration of the power consumption of the DC is required at all time. In the present approach, we neglect Fridgen et al. (2017)'s major drawback of the mandatory utilization of a balancing mechanism. We show that, under certain circumstances, two market areas can exchange BP without using a balancing mechanism. This is due to the fact that BP can be compensated by the two BP markets when a positive BP call in one BP market location and a negative BP call in the second BP market location has to be satisfied at the same time. The compensation is based on the fact that energy systems at different locations are in different and thus possibly contradictory situations due to different influencing factors (e.g., weather and daytime). Thus, virtually interconnecting these energy systems by spatial load migration offers a decisive advantage over temporary flexibility since it allows to exploit these potential differences.

The virtual interconnection is achieved by coupling two different types of flows, namely the power flow, which is typically “transmitted” by power grids, and the information flow, which is typically “transmitted” by information networks. Thus, the objective of the model is to intelligently control information flows in

order to manage power flows, which are altered to meet the precise demands of the selected BP markets. The process involves the coupling of two different sectors, and, as such, our project extends the existing definition of sector coupling, introducing a new, highly relevant component, the information sector, which should always be considered whenever discussing potential concepts of cross-sector flexibility.

By virtually interconnecting distant BP markets, our approach forges new links in the worldwide power system. In our exemplary evaluation scenario, we interconnected two different markets in two different countries, as well as a third balancing mechanism in a third country. As a result, we were able to illustrate two different effects: Firstly, there were periods in which we migrated the balancing potential of the balancing mechanism to a distant power grid. Secondly, and even more importantly, we identified numerous periods in which it was possible to use opposite balancing demands, resulting in the simultaneous stabilization of distant power grids without the use of a conventional balancing mechanism. Based on this second effect, and provided that our model is deployed to multiple BP markets around the world, this approach makes it possible to reduce the global demand for conventional balancing mechanisms.

As introduced, in near future, new, omnipresent trends in information technology, such as the Internet of Things (Atzori et al., 2010) and blockchain (Schweizer et al., 2017; Swan, 2015), might further increase demand for computational power. This trend of decentralization in provision of computational power will continue to result in new approaches, such as the current concept of fog computing (Bonomi et al., 2012; Vaquero and Rodero-Merino, 2014). The idea of fog computing is to further decentralize the power-intensive processing of data by sharing the workload between connected, distributed devices with computational power. This concept aims to increase the use of the existing global computational capacity, e.g., of residential personal computers, or even small devices such as Internet of Things devices (e.g., smart watches). This development of increasing demand for computing power and the decentralization of computational power provision provide ground to useful interconnect more different BP markets to enable a compensation of BP demand. In this case, for every incoming computational request, decisions about the amount of computational capacity to be drawn from decentralized devices could also take into consideration the power grid's balancing demand. This would provide the opportunity for the fine-grained

adjustment of spatial energy demand around the globe. Therefore, in our opinion, the information sector is a key part of the digital power system and will play an increasingly important role in its future.

The increasing share of volatile (decentralized) energy generation based on renewable energy sources (e.g., wind and solar energy) comes with a high demand for power transmission capacity, which is necessary to continually safeguard the balance between power generation and power demand. However, in today's society, extensions to the power transmission capacity are met with resistance due to high costs and transmission losses (Fridgen et al., 2017). Consequently, by coupling the information and power sectors, which are both becoming increasingly decentralized, we demonstrate how information flows could be used as a substitute for energy flows, and might thus reduce the need for power grid expansion. This alternative could be geographically replicated around the world without transmission losses, based on an existing decentralized IT infrastructure. Thus, our approach contributes to the world-wide connection of distributed power grids without the need for additional transmission lines. Given that this change is enabled by the information sector, this sector should not be neglected when considering the development of tomorrow's power systems. Accordingly, operators of information and communication infrastructure might be involved in the discourse over how to ensure grid stability given the increasing share of renewables. In short, the model we have introduced here requires a cross-sectional dialogue and transdisciplinary collaboration, which we want to motivate with this article.

This cooperation should not only involve practitioners, but also, in particular, political decision-makers. Policy-makers are responsible for adjusting the regulatory framework in order to allow fair and sustainable cooperation between the two sectors. This includes, *inter alia*, a discussion about opening up balancing markets to new, previously excluded players (e.g., from the information and communication sector), for whom the existing requirements are not achievable. Moreover, the future digital energy system could profit from the standardization of both the information and power sectors. Standardization in the information sector would establish the interchangeability of IT infrastructure in as many locations as possible. Standardization in the power sector would provide barrier-free and consistent options to market flexibility in unified market models.

5.2 Directions for further research

Our model of an exemplary setup for load migration suggests different opportunities for further research. First of all, we have made some simple assumptions in order to provide an initial estimate of the economic potential. In our future research, we want to address these assumptions in order to improve our model and apply it to a wider range of problems. Moreover, we have not yet gone into detail about how the regulatory framework should be designed in order to effectively incentivize cooperation between the power sector and the information sector.

The extension of the setup components represents another possible direction for further, interdisciplinary research. Our model is based on large-scale, geographically distributed DCs. However, as illustrated above, we also observe a trend towards fine-grained, distributed IT resources (i.e., fog computing or regular data centers). The pooling of these resources – as well as problems associated with the pooling, e.g., regarding IT security or coordination of IT devices – provides new opportunities for future research in different disciplines. Last but not least, we focus on existing national BP markets. However, we observe a trend towards innovative regional market places or even bilateral contracts for spatial load migration. Thus, the design of innovative market mechanisms for providing flexibility among consumers in the information sector offers another possibility for further research.

5.3 Limitations and conclusion

In this paper, we presented a model which enables the virtual interconnection of distant BP markets. This model is based on a setup consisting of at least three DCs in different locations. Two of the DCs participate in the local BP markets and the third DC has access to a balancing mechanism. To deliver BP with this setup, incoming requests are intelligently dispatched to the three DCs in order to control their workload and, thus, control their power demand.

To allow a first estimation of our model's potential, we simulated the interconnection of the German and the Dutch BP markets. The third DC in this setup is located in Finland and is able to access a local balancing mechanism. This simulated study is largely based on real-world data. The results of our evaluation indicate an economic incentive to apply our model. Moreover, the evaluation indicates that our model allows the compensation of BP demand without the need for a conventional balancing mechanism. This compensation

is based on the fact that, in several periods, a positive BP call in one of the BP market locations (Germany or the Netherlands) meets a negative BP call in the second BP market location. Finally, we discussed policy implications and new trends in technology which increase the importance of our model.

However, when interpreting the results of our paper it is important to consider the limitations. First of all, we demonstrate the economic potential of our model, but do not investigate how this economic potential should be allocated among the stakeholders in order to incentivize their participation. Secondly, our model is based on three major assumptions which we describe in the article. Although the provided approach to compensate BP demand would also be applicable without these assumptions, the economic potential may differ from the results we have presented. Lastly, we investigate the economic potential for a specific setting which is based on real-world data from existing BP markets. Accordingly, the economic potential we identified is not necessarily present in all other BP markets around the world. Furthermore, we implicitly assume that the behavior of the DCs does not have an impact on the price level for BP, i.e., the DCs are price-takers. However, the large-scale use of our model will probably have an impact on BP market price levels and, thus, on the economic potential of the model.

In this article, we present an example of an innovative model which virtually interconnects different markets using DCs. Accordingly, we couple two different sectors; namely the power and information sectors. New trends in these sectors moderating the boundaries between the two and contribute to the overarching concept of a future digital energy system. By illustrating a way to connect the power and information sectors, we contribute to this development and provide a new approach to shape the future energy system.

References

- Adnan, M.A., Sugihara, R., Gupta, R.K., 2012. Energy Efficient Geographical Load Balancing via Dynamic Deferral of Workload, in: 2012 IEEE Fifth International Conference on Cloud Computing. IEEE, pp. 188–195. <https://doi.org/10.1109/CLOUD.2012.45>
- Alphabet Inc., 2017. Efficiency: How we do it – Data Centers [WWW Document]. URL <https://www.google.com/about/datacenters/efficiency/internal/> (accessed 2.2.17).
- Andrzejak, A., Kondo, D., Yi, S., 2010. Decision Model for Cloud Computing under SLA Constraints.
- Atzori, L., Iera, A., Morabito, G., 2010. The Internet of Things: A survey. *Comput. Networks* 54, 2787–2805. <https://doi.org/10.1016/J.COMNET.2010.05.010>
- Avelar, V., Azevedo, D., French, A., 2012. PUE™: A Comprehensive Examination of the Metric [WWW Document]. URL [https://datacenters.lbl.gov/sites/all/files/WP49-PUE A Comprehensive Examination of the Metric_v6.pdf](https://datacenters.lbl.gov/sites/all/files/WP49-PUE_A_Comprehensive_Examination_of_the_Metric_v6.pdf) (accessed 2.2.17).
- Bajpai, P., Singh, S.N., 2007. Fuzzy Adaptive Particle Swarm Optimization for Bidding Strategy in Uniform Price Spot Market. *IEEE Trans. Power Syst.* 22, 2152–2160. <https://doi.org/10.1109/TPWRS.2007.907445>
- Bonomi, F., Milito, R., Zhu, J., Addepalli, S., 2012. Fog computing and its role in the internet of things, in: Proceedings of the First Edition of the MCC Workshop on Mobile Cloud Computing. ACM Press, New York, pp. 13–16. <https://doi.org/10.1145/2342509.2342513>
- Buijs, P., Bekaert, D., Cole, S., Van Hertem, D., Belmans, R., 2011. Transmission investment problems in Europe: Going beyond standard solutions. *Energy Policy* 39, 1794–1801. <https://doi.org/10.1016/j.enpol.2011.01.012>
- CAISO, 2015. Fifth Replacement Electronic Tariff [WWW Document]. URL https://www.caiso.com/Documents/ConformedTariff_Jun3_2015.pdf (accessed 2.2.17).
- Chatzivasilieiadis, S., Ernst, D., Andersson, G., 2013. The Global Grid. *Renew. Energy* 57, 372–383. <https://doi.org/10.1016/J.RENENE.2013.01.032>
- Consentec, 2014. Beschreibung von Regelleistungskonzepten und Regelleistungsmarkt [WWW Document]. URL http://www.consentec.de/wp-content/uploads/2014/08/Consentec_50Hertz_Regelleistungsmarkt_de_201402271.pdf (accessed 2.2.17).
- Derakhshan, G., Shayanfar, H.A., Kazemi, A., 2016. The optimization of demand response programs in smart grids. *Energy Policy* 94, 295–306. <https://doi.org/10.1016/j.enpol.2016.04.009>
- Ehrlich, L.G., Klamka, J., Wolf, A., 2015. The potential of decentralized power-to-heat as a flexibility option for the german electricity system: A microeconomic perspective. *Energy Policy* 87, 417–428. <https://doi.org/10.1016/j.enpol.2015.09.032>
- ENTSO-E, 2018. ENTSO-E Transparency Platform [WWW Document]. URL <https://transparency.entsoe.eu/> (accessed 9.3.18).
- Feuerriegel, S., Neumann, D., 2014. Measuring the financial impact of demand response for electricity retailers. *Energy Policy* 65, 359–368. <https://doi.org/10.1016/j.enpol.2013.10.012>
- Fingrid, 2017. Fingrid - Balancing power market [WWW Document]. URL [http://www.fingrid.fi/en/customers/Balance services/management/powermarket/Pages/default.aspx](http://www.fingrid.fi/en/customers/Balance_services/management/powermarket/Pages/default.aspx) (accessed 2.2.17).
- Flinkerbusch, K., Heuterkes, M., 2010. Cost reduction potentials in the German market for balancing power. *Energy Policy* 38, 4712–4718. <https://doi.org/10.1016/j.enpol.2010.04.038>
- Fridgen, G., Häfner, L., König, C., Sachs, T., 2016. Providing Utility to Utilities: The Value of Information Systems Enabled Flexibility in Electricity Consumption. *J. Assoc. Inf. Syst.* 17, 537–563.
- Fridgen, G., Häfner, L., König, C., Sachs, T., 2015. The Value of IS-Enabled Flexibility in Electricity Demand

- a Real Options Approach, in: The International Conference on Wirtschaftsinformatik. Osnabrück, Germany.
- Fridgen, G., Keller, R., Thimmel, M., Wederhake, L., 2017. Shifting load through space—The economics of spatial demand side management using distributed data centers. *Energy Policy* 109, 400–413. <https://doi.org/10.1016/J.ENPOL.2017.07.018>
- Fridgen, G., Mette, P., Thimmel, M., 2014. The Value of Information Exchange in Electric Vehicle Charging, in: 35th International Conference on Information Systems (ICIS). Auckland, New Zealand.
- Gandhi, A., Gupta, V., Harchol-Balter, M., Kozuch, M.A., 2010. Optimality analysis of energy-performance trade-off for server farm management. *Perform. Eval.* 67, 1155–1171. <https://doi.org/10.1016/j.peva.2010.08.009>
- Ganesh, L., Weatherspoon, H., Marian, T., Birman, K., 2013. Integrated Approach to Data Center Power Management. *IEEE Trans. Comput.* 62, 1086–1096. <https://doi.org/10.1109/TC.2013.32>
- Gellings, C.W., 2015. Let's Build a Global Power Grid [WWW Document]. *IEEE Spectr.* URL <https://spectrum.ieee.org/energy/the-smarter-grid/lets-build-a-global-power-grid> (accessed 9.4.18).
- Ghatikar, G., Ganti, V., Matson, N., Piette, M.A., 2012. Demand response opportunities and enabling technologies for data centers: Findings from field studies [WWW Document].
- Ghatikar, G., Mary, A.P., Fujita, S., McKane, A., Dudley, J., Radspieler, A., 2010. Demand Response and Open Automated Demand Response Opportunities for Data Centers [WWW Document].
- Gmach, D., Rolia, J., Bash, C., Chen, Y., Christian, T., Shah, A., Sharma, R., Wang, Z., 2010. Capacity planning and power management to exploit sustainable energy, in: 2010 International Conference on Network and Service Management. *IEEE*, pp. 96–103. <https://doi.org/10.1109/CNSM.2010.5691329>
- Gmach, D., Rolia, J., Cherkasova, L., Kemper, A., 2007. Workload Analysis and Demand Prediction of Enterprise Data Center Applications, in: 2007 IEEE 10th International Symposium on Workload Characterization. *IEEE*, pp. 171–180. <https://doi.org/10.1109/IISWC.2007.4362193>
- Goddard, G., Klose, J., Backhaus, S., 2014. Model development and identification for fast demand response in commercial HVAC systems. *IEEE Trans. Smart Grid* 5, 2084–2092. <https://doi.org/10.1109/TSG.2014.2312430>
- Grein, A., Pehnt, M., 2011. Load management for refrigeration systems: Potentials and barriers. *Energy Policy* 39, 5598–5608. <https://doi.org/10.1016/j.enpol.2011.04.040>
- Hintemann, R., Clausen, J., 2014. Rechenzentren in Deutschland: Eine Studie zur Darstellung der wirtschaftlichen Bedeutung und der Wettbewerbssituation. Berlin.
- Jang, D., Eom, J., Jae Park, M., Jeung Rho, J., 2016. Variability of electricity load patterns and its effect on demand response: A critical peak pricing experiment on Korean commercial and industrial customers. *Energy Policy* 88, 11–26. <https://doi.org/10.1016/j.enpol.2015.09.029>
- Jheng, J.J., Tseng, F.H., Chao, H.C., Chou, L. Der, 2014. A novel VM workload prediction using Grey Forecasting model in cloud data center, in: International Conference on Information Networking. *IEEE*, pp. 40–45. <https://doi.org/10.1109/ICOIN.2014.6799662>
- Keller, R., Häfner, L., Sachs, T., Fridgen, G., 2019. Scheduling Flexible Demand in Cloud Computing Spot Markets. *Bus. Inf. Syst. Eng.* <https://doi.org/10.1007/s12599-019-00592-5>
- Keller, R., König, C., 2014. A Reference Model to Support Risk Identification in Cloud Networks, in: 35th International Conference on Information Systems (ICIS). Auckland, New Zealand.
- Kishore, T.S., Singal, S.K., 2014. Optimal economic planning of power transmission lines: A review. *Renew. Sustain. Energy Rev.* 39, 949–974. <https://doi.org/10.1016/j.rser.2014.07.125>
- Kong, F., Liu, X., 2014. A Survey on Green-Energy-Aware Power Management for Datacenters. *ACM Comput. Surv.* 47, 1–38. <https://doi.org/10.1145/2642708>

- Koomey, J., Taylor, J., 2015. New data supports finding that 30 percent of servers are “Comatose”, indicating that nearly a third of capital in enterprise data centers is wasted.
- Krcmar, H., 2015. Informationsmanagement. Springer, Berlin, Heidelberg.
- Krishna, V., 2010. Auction Theory, Second Edition. Academic press.
- Lin, M., Wierman, A., Thereska, E., 2013. Dynamic Right-Sizing for Power-Proportional Data Centers. *IEEE/ACM Trans. Netw.* 21, 1378–1391. <https://doi.org/10.1109/TNET.2012.2226216>
- Lujano-Rojas, J.M., Monteiro, C., Dufo-López, R., Bernal-Agustín, J.L., 2012. Optimum residential load management strategy for real time pricing (RTP) demand response programs. *Energy Policy* 45, 671–679. <https://doi.org/10.1016/j.enpol.2012.03.019>
- Lütticke, M., 2017. Opponents strike back at planned power line [WWW Document]. URL <http://www.dw.de/opponents-strike-back-at-planned-power-line/a-17410969> (accessed 2.2.17).
- Mao, M., Humphrey, M., 2012. A Performance Study on the VM Startup Time in the Cloud, in: 2012 IEEE Fifth International Conference on Cloud Computing. IEEE, pp. 423–430. <https://doi.org/10.1109/CLOUD.2012.103>
- Meisner, D., Gold, B.T., Wenisch, T.F., Meisner, D., Gold, B.T., Wenisch, T.F., Meisner, D., Gold, B.T., Wenisch, T.F., 2009. PowerNap, in: Proceeding of the 14th International Conference on Architectural Support for Programming Languages and Operating Systems - ASPLOS '09. ACM Press, New York, New York, USA, p. 205. <https://doi.org/10.1145/1508244.1508269>
- Müller, G., Rammerstorfer, M., 2008. A theoretical analysis of procurement auctions for tertiary control in Germany. *Energy Policy* 36, 2620–2627. <https://doi.org/10.1016/j.enpol.2008.03.017>
- Nord Pool Spot, 2018. Historical Market Data [WWW Document]. URL <http://www.nordpoolspot.com/historical-market-data/> (accessed 9.3.18).
- Palensky, P., Dietrich, D., 2011. Demand Side Management: Demand Response, Intelligent Energy Systems, and Smart Loads. *IEEE Trans. Ind. Informatics* 7, 381–388. <https://doi.org/10.1109/TII.2011.2158841>
- Papagiannis, G., Dagoumas, A., Lettas, N., Dokopoulos, P., 2008. Economic and environmental impacts from the implementation of an intelligent demand side management system at the European level. *Energy Policy* 36, 163–180. <https://doi.org/10.1016/j.enpol.2007.09.005>
- Qu, C., Calheiros, R.N., Buyya, R., 2017. Mitigating impact of short-term overload on multi-cloud web applications through geographical load balancing. *Concurr. Comput. Pract. Exp.* 29, e4126. <https://doi.org/10.1002/cpe.4126>
- Rammerstorfer, M., Wagner, C., 2009. Reforming minute reserve policy in Germany: A step towards efficient markets? *Energy Policy* 37, 3513–3519. <https://doi.org/10.1016/j.enpol.2009.03.056>
- Rasmussen, M.G., Andresen, G.B., Greiner, M., 2012. Storage and balancing synergies in a fully or highly renewable pan-European power system. *Energy Policy* 51, 642–651. <https://doi.org/10.1016/J.ENPOL.2012.09.009>
- Rebours, Y.G., Kirschen, D.S., Trotignon, M., Rossignol, S., 2007a. A Survey of Frequency and Voltage Control Ancillary Services—Part I: Technical Features. *IEEE Trans. Power Syst.* 22, 350–357. <https://doi.org/10.1109/TPWRS.2006.888963>
- regelleistung.net, 2018. Internetplattform zur Vergabe von Regelleistung [WWW Document]. URL <https://www.regelleistung.net/ip/action/index> (accessed 8.31.18).
- Schott, P., Ahrens, R., Bauer, D., Hering, F., Keller, R., Pullmann, J., Schel, D., Schimmelpfennig, J., Simon, P., Weber, T., Abele, E., Bauernhansl, T., Fridgen, G., Jarke, M., Reinhart, G., 2018. Flexible IT platform for synchronizing energy demands with volatile markets. *it - Inf. Technol.* 60, 155–164. <https://doi.org/10.1515/itit-2018-0001>
- Schweizer, A., Schlatt, V., Urbach, N., Fridgen, G., 2017. Unchaining Social Businesses – Blockchain as the Basic Technology of a Crowdfunding Platform, in: 38th International Conference on Information

Systems (ICIS). Seoul.

- Shehabi, A., Smith, S., Sartor, D., Brown, R., Herrlin, M., Koomey, J., Masanet, E., Horner, N., Azevedo, I., Lintner, W., 2016. United States Data Center Energy Usage Report. Lawrence Berkeley Natl. Lab. Berkeley, California. LBNL-1005775.
- Simon, H.A., 1979. Rational Decision Making in Business Organizations. *Am. Econ. Rev.* 69, 493–513.
- Swan, M., 2015. *Blockchain : blueprint for a new economy*. O'Reilly Media.
- Swider, D.J., 2006. *Handel an Regelenergie- und Spotmärkten: Methoden zur Entscheidungsunterstützung für Netz- und Kraftwerksbetreiber*. GWV Fachverlag GmbH, Wiesbaden.
- TenneT, 2018. Data export [WWW Document]. URL http://www.tennet.org/english/operational_management/export_data.aspx (accessed 9.3.18).
- TenneT, 2012. Manual Bidding of Regulating and Reserve Power (RRP) [WWW Document]. URL www.tennet.org. (accessed 7.10.18).
- Tolia, N., Wang, Z., Marwah, M., Bash, C., Ranganathan, P., Zhu, X., 2008. Delivering energy proportionality with non energy-proportional systems - optimizing the ensemble. *HotPower*.
- Tseng, F.-H., Wang, X., Chou, L.-D., Chao, H.-C., Leung, V.C.M., 2018. Dynamic Resource Prediction and Allocation for Cloud Data Center Using the Multiobjective Genetic Algorithm. *IEEE Syst. J.* 12, 1688–1699. <https://doi.org/10.1109/JSYST.2017.2722476>
- Van Hulle, F., Fichaux, N., Sinner, A.F., Morthost, P.E., Munksgaard, J., Ray, S., 2010. Powering Europe: Wind Energy and the Electrical Grid [WWW Document]. URL http://www.ewea.org/fileadmin/files/library/publications/reports/Grids_Report_2010.pdf (accessed 2.2.17).
- Vandezande, L., Meeus, L., Belmans, R., Saguan, M., Glachant, J.-M., 2010. Well-functioning balancing markets: A prerequisite for wind power integration. *Energy Policy* 38, 3146–3154. <https://doi.org/10.1016/j.enpol.2009.07.034>
- Vaquero, L.M., Rodero-Merino, L., 2014. Finding your Way in the Fog: Towards a Comprehensive Definition of Fog Computing. *ACM SIGCOMM Comput. Commun. Rev.* 44, 27–32. <https://doi.org/10.1145/2677046.2677052>
- Vennemann, P., Gruber, K.H., Haaheim, J.U., Kunsch, A., 2011. Pumped Storage Plants – Status and Perspectives. *VGB PowerTech* 91, 32–28.
- Vieira, C.C.A., Bittencourt, L.F., Madeira, E.R.M., 2015. A Scheduling Strategy Based on Redundancy of Service Requests on IaaS Providers, in: 2015 23rd Euromicro International Conference on Parallel, Distributed, and Network-Based Processing. IEEE, pp. 497–504. <https://doi.org/10.1109/PDP.2015.80>
- Wang, J., Zhong, H., Tang, W., Rajagopal, R., Xia, Q., Kang, C., Wang, Y., 2017. Optimal bidding strategy for microgrids in joint energy and ancillary service markets considering flexible ramping products. *Appl. Energy* 205, 294–303. <https://doi.org/10.1016/J.APENERGY.2017.07.047>
- Whitney, J., Delforge, P., 2014. Data Center Efficiency Assessment [WWW Document]. URL <https://www.nrdc.org/energy/files/data-center-efficiency-assessment-IP.pdf> (accessed 2.2.17).
- Wierman, A., Liu, Z., Liu, I., Mohsenian-Rad, H., 2014. Opportunities and challenges for data center demand response, in: International Green Computing Conference. IEEE, pp. 1–10. <https://doi.org/10.1109/IGCC.2014.7039172>
- Zhang, Y., Wang, Y., Wang, X., 2012. Electricity Bill Capping for Cloud-Scale Data Centers that Impact the Power Markets, in: International Conference on Parallel Processing. IEEE, pp. 440–449. <https://doi.org/10.1109/ICPP.2012.23>
- Zhou, M., Zhang, R., Xie, W., Qian, W., Zhou, A., 2010. Security and Privacy in Cloud Computing: A Survey, in: 2010 Sixth International Conference on Semantics, Knowledge and Grids. IEEE, pp. 105–112. <https://doi.org/10.1109/SKG.2010.19>

Zissis, D., Lekkas, D., 2012. Addressing cloud computing security issues. *Futur. Gener. Comput. Syst.* 28, 583–592. <https://doi.org/10.1016/j.future.2010.12.006>