Energy Policy 109 (2017) 400–413

Contents lists available at ScienceDirect

Energy Policy

journal homepage: www.elsevier.com/locate/enpol

Shifting load through space–The economics of spatial demand side management using distributed data centers

Gilbert Fridgen[a], Robert Keller[a,*], Markus Thimmel[b], Lars Wederhake[b]

[a] *FIM Research Center, University of Bayreuth, Wittelsbacherring 10, 95444 Bayreuth, Germany*
[b] *FIM Research Center, University of Augsburg, Universitaetsstrasse 12, 86159 Augsburg, Germany*

ARTICLE INFO

Keywords:
Load migration
Economic potential
Balancing power
Simulation-based case study
Demand response
Demand-side management

ABSTRACT

Demand-side flexibility (DSF) in the electricity grid has become an active research area in recent years. While temporal flexibility (e.g. load shedding, load shifting) is already discussed intensively in literature, spatial load migration still is an under-researched type of DSF. Spatial load migration allows us to instantly migrate power-consuming activities among different locations. Data centers (DCs) are power-intensive and process information goods. Since information goods are easily transferable through communication networks, power-intensive processing of information goods is not necessarily tied to a specific location. Consequently, geographically distributed DCs inherit—in theory—a considerable potential to globally migrate load. We analyze the economics of spatially migrating load to provide balancing power using geographically distributed DCs. We assure that neither of the participating electricity grids will be burdened by this mechanism. By using historical data to evaluate our model, we find reasonable economic incentives to migrate positive as well as negative balancing power. In addition, we find that current scenarios favor the migration of negative balancing power. Our research thus reveals realistic opportunities to virtually transfer balancing power between different market areas worldwide.

1. Motivation

In power markets, supply and demand must be so tightly coupled that they line up all times (Müller and Rammerstorfer, 2008; Rammerstorfer and Wagner, 2009). However, it is difficult to know a priori what the exact demand for power will be (Flinkerbusch and Heuterkes, 2010). Considering the ever-growing power generation from intermittent renewables, the supply side introduces even more uncertainty (Vandezande et al., 2010). The resulting high levels of volatility demand elevated levels of flexibility (Ehrlich et al., 2015; Strbac, 2008). Flexibility is the potential to balance deviations from the scheduled power generation or demand caused by prediction errors (Eurelectric, 2014). The “intelligent control” (Buhl and Jetter, 2009) of demand-side resources by information and communication technologies (ICT) increases a power grid's ability to react on higher levels of volatility (Fridgen et al., 2015; Strbac, 2008). Unsurprisingly, examining the potential of demand-side flexibility (DSF) has become an active field of research in recent years, e.g. electric vehicles (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 commercial and industrial processes (Jang et al., 2016). As these examples illustrate, the vast majority of today's approaches to providing flexibility are variants of temporal flexibility (load shifting and load shedding).

Another approach to providing flexibility is to spatially migrate load. The exchange is favorable for both importers and exporters of power for two reasons: first, excess power in one location can neutralize a deficit in another, and second, some markets can provision flexibility more cost-efficiently and/or in a less carbon-emitting manner than others (Van Hulle et al., 2010; Vennemann et al., 2011). Elements for provisioning flexibility efficiently are the mix of energy generation, efficient storage facilities, and the potential for adjusting load. The potential of the first two of these components, however, is heavily influenced by geographical realities (Vennemann et al., 2011). For this reason, it is imperative to interconnect markets to reap the benefits from reducing market inefficiencies. Nevertheless, many power line construction projects fail because of excessive initial costs (Kishore and Singal, 2014), insecure return on investment (Buijs et al., 2011), protests by local citizens (Lütticke, 2017), and high transmission costs (Vennemann et al., 2011).

---

[*] Corresponding author.

E-mail addresses: gilbert.fridgen@fim-rc.de (G. Fridgen), robert.keller@fim-rc.de (R. Keller), markus.thimmel@fim-rc.de (M. Thimmel), lars.wederhake@fim-rc.de (L. Wederhake).

http://dx.doi.org/10.1016/j.enpol.2017.07.018
Received 4 March 2017; Received in revised form 6 July 2017; Accepted 7 July 2017
0301-4215/

An alternative concept to spatially migrating load is to instantly migrate power-consuming activities between different locations. Processing and delivering information goods are very power-consuming activities. This is because these activities are, in theory perfectly location independent (Krcmar, 2015). Data centers (DCs) processing and delivering information goods are very power-intensive. As an example, all US-based DCs contribute some 2% to the country's total electricity consumption (Koomey, 2011). Accordingly, a setting consisting of geographically distributed DCs could enable the spatial migration of load over long distances, relying on communication networks instead of the power grid.

Migrating load requires the participation of economic entities. The actions of economic entities are mainly driven by economic rationale (Simon, 1979). Therefore, economic entities will only adopt the migration of load if it is economically feasible. Since the intentional in-/decrease of load is one possible source of flexibility, geographically distributed DCs can deliver flexibility in one location by migrating load from this location to a remote location. However, this might also result in unintentional in-/decreases of load in the remote location. Consequently, a major challenge of spatially migrating load is to avoid one power grid improving grid stability at the expense of another. This could result from the additional power imbalances that might be introduced.

> *Thus, the objective of this paper is to analyze the economic feasibility (cash flow from operating activities) of spatially migrating load in order to provide flexibility, burdening neither of the participating power grids by potentially introduced additional power imbalances. We demonstrate the economic feasibility of spatially migrating load enabled by geographically distributed DCs.*

2. Demand-side flexibility approaches

Despite recent research efforts examining the potential of DSF, according to the Energy Policy's manuscript by Feuerriegel and Neumann (2014) “little is known about the economic dimension [of DSF] and further effort is strongly needed to realistically quantify the financial impact”. Because there are different approaches to DSF, economic analyses must take care of the decisive differences.

DSF is generally considered to be based on two types of approaches – load shedding and load shifting (Derakhshan et al., 2016; Feuerriegel and Neumann, 2014). Load shifting refers to the concept of postponing or putting forward an energy-consuming activity in order to reduce load during peak hours – e.g. a charging process of an electric vehicle (Fridgen et al., 2016). Load shedding refers to the concept of ceasing or not starting a planned energy-consuming activity without resuming it later in time – e.g. switching off street lights (Papagiannis et al., 2008).

In the domain of DC management, the two types of DSF approaches that are generally applicable are extended by the idea of spatially “shifting” load. This is done by assigning requests to geographically distributed DCs. This idea is commonly referred to as load migration (e.g. Adnan et al., 2012; Wierman et al., 2014). In the section below, we briefly describe applications of load shifting and load shedding in the context of DCs, and then provide information about related works with regard to load migration.

2.1. Load shifting & load shedding

Load shifting requires load to be time-flexible. In the domain of DCs, some types of requests are inflexible but some are delay-tolerant (Gmach et al., 2010). The latter type comprises scientific computing, routine tasks such as batch processing, and more recently, bitcoin mining (Lewenberg et al., 2015; e.g. Nakamoto, 2008). Some DCs solely serve a single type of request. If this type is delay-tolerant, the load shifting potential can be as high as the difference between maximum load and the load at idle state.

The possibility of creating DSF through load shedding typically comes with quality degradation under the constraints as quality-of-service (QoS) requirements and service level agreements (SLAs) as outlined by Wierman et al. (2014). An example of this is the growing number of big data algorithms as they are used for targeting ads. A DC can reduce consumption by targeting ads less effectively (e.g. Baek and Chilimbi, 2010). Empirical studies by the Lawrence Berkeley National Laboratory found that 5% of the load can typically be shed in 5 min and 10% of the load can be shed in 15 min; and that these can be achieved without changes to how the IT workload is handled (Ghatikar et al., 2012, 2010).

2.2. Load migration

Geographically distributed DCs provide the opportunity of migrating load between locations of the DCs (Ghatikar et al., 2012, 2010; Kong and Liu, 2014) and thus contribute to grid stability by intelligently assigning (dispatching) the incoming workload to geographically distributed DCs.

There are a few contributions conducting economic analyses of load migration by DCs according to power price differences (e.g. Qureshi et al., 2009; Li et al., 2012; Zhang et al., 2012; Camacho et al., 2014). These approaches indirectly contribute to medium- and long-term power grid stability, since the power market prices are, in general, determined by the available power supply and demand. However, price signals on power markets are not suitable for helping on short-term grid stability for two reasons: first, apart from perhaps real-time markets, lead time is too long, i.e. the time between gate closure and the time the contract becomes valid. Second, when trading on several markets, there can be no guarantee that the DC's primary objective of reducing procurement costs will be aligned with the grid's stability objective. This is because absolute price differences between markets can outplay relative (intra-market) price differences. So, the DC might benefit from migrating load to another market although its relative price is far above average. Chiu et al. (2012) propose a concept for grid balancing by intelligently dispatching incoming workload between geographically distributed DCs pursuant to local real-time price signals, assuming they exist. However, in the majority of markets they do not.

The listed contributions illustrate that there are few options to trade flexibility. To the best of our knowledge, previous research focuses on doing so on power markets and innovative (barely existing) bilateral products, only. Müller and Rammerstorfer (2008) show that delivering balancing power (BP) is not only another and already existing option for trade flexibility but is specifically designed for this purpose. Rebours et al. (2007a) illustrate that trades of BP are usually placed on so-called BP markets. In contrast to the referred articles, this paper provides flexibility by spatially migrating load and markets this flexibility on a BP market. We refer to this process as BP migration.

There is a potential major drawback of load migration, however: the challenge of how to avoid one power grid improving its stability at the expense on another. For this reason, we base our analysis on the collaboration of balancing mechanisms (BM), e.g. pumped hydropower plants, and geographically distributed DCs. In the following section, we describe a process based on that collaboration allowing BP migration without the potential drawback.

3. BP migration process

3.1. Setup

Generally, it is possible to apply the process presented in this section to a bidirectional migration of BP (between two locations). Because, in this paper, we merely strive to demonstrate the economic feasibility, we choose an irreducible setup considering a unidirectional BP migration. Mirroring the setup to the opposite direction facilitates