2020 International Conferences on Internet of Things (iThings) and IEEE Green Computing and Communications (GreenCom) and IEEE Cyber, Physical and Social Computing (CPSCom) and IEEE Smart Data (SmartData) and IEEE Congress on

Decentral load control for data centers

Felix Uster, Franz Plocksties, Dirk Timmermann
Institute of Applied Microelectronics and Computer Engineering
University of Rostock
Rostock, Germany
felix.uster@uni-rostock.de, franz.plocksties@uni-rostock.de , dirk.timmermann@uni-rostock.de

***Abstract*—The regulation and stabilization of the power grid requires various system services that have been provided so far primarily by conventional power plants. With the future shutdown of coal-fired and nuclear power plants, investigations into the provision of these services will become all the more important in order to ensure the safe operation of the grid. Currently, photovoltaic and wind can only provide limited so-called control power, so that new methods such as the regulation of large consumers will become interesting. Adapting the electricity requirements of consumers to the current grid situation has a comparable effect to the provision of additional balancing power on the generation side. Some consumers like data centers can provide primary control power at high speed. In this paper we present how data centers can autonomously support grid stabilization with minimal impact on the daily operation. Additional benefits include an higher efficiency and the ability to take advantage of any future financial bonuses by regulation authorities or electricity suppliers.**

***Index Terms*—decentralized grid stabilization, data center, load control, demand response, grid adaptation**

I. Introduction

In recent years, the growth in data center power consumption stalled, but it still accounts for a significant share of total power consumption and new applications such as cloud-gaming and machine learning are emerging. Highest share of the power consumption arises from the processors and the necessary cooling equipment. In the USA, the National Institute of Standards and Technology and the Department of Energy registered *demand response*, the adaptation of power consumption under certain circumstances as considerable component on the consumer side for future smart grids [1], [2]. In general, it is assumed that electricity demand is inflexible. However, a grid-friendly behaviour requires the ability to adapt the demand as quickly as possible whereby the specific operation adaptation depends on the type of consumer and its characteristics. Data centers are often referred to as flexible loads because of the non time-critical workload that can be shifted within boundaries [3], [4]. *Demand response* includes a signal by the grid operator to decrease the power consumption during high demand or high electricity price intervals. Our approach, called *decentral load control*, operates independently from any regulator and is entirely focused on the stabilization of the power grid. In our previous research we already proved the positive effect of *decentral load control* enabled consumers on the *synchronous grid of Continental Europe* [5].

The novelty of this work is the investigation of one specific way to implement *decentral load control* in and its effect on data centers and the stabilizing effect for the power grid. In [5] we presented general approaches for load control independent of the load and already mentioned that these mechanisms should be optimized for specific applications. Therefore, additional research is necessary to conclude how to adjust the power consumption and how it affects the operation of a load, in this work for the application in data

978-1-7281-7647-5/20/$31.00 ©2020 IEEE

DOI 10.1109/iThings-GreenCom-CPSCom-SmartData-Cybermatics50389.2020.00070

centers, and minimize the effect of this additional control for an autonomous grid support. Within data centers different approaches are possible to adjust the power consumption, discussed in the next section, and we focus on the widely available and simple implementable way by utilizing basic modern processor capabilities.

In 2008, Europe launched the *Code of Conduct for Energy Efficiency in Data Centers* as response to the increasing power consumption of data centers to reduce its impact on environment, economy, and energy supply [6]. Regarding the future grid situation heavily based on renewable energies, we assume that the presented decentral load control approach seems appropriate as an extension for the *Code of Conduct for Energy Efficiency in Data Centers* as it also supports the integration of renewable energy into the power grid. Conventional power plants provide inertia and control power, both necessary for a stable grid [7]. But renewable energies provide hardly any inertia and cannot reliably provide control power, increasing the demand for alternative concepts like *decentral load control*. Comparable to the increasingly decentral electricity generation, our approach distributes control and stabilization functions throughout the power grid.

This paper is organized as follows: Section II primarily describes the methods used for data center's power consumption manipulation from previous work compared to our conception. Section III compares the already known *demand response* approach with our idea of the decentralized use of data centers for power grid stabilization. The section *DVFS for Grid Stabilization* (Dynamic Voltage and Frequency Scaling - DVFS) explains how we utilize basic processor principles to adjust data center's power consumption for a grid supporting effect under the necessary condition that the operation should not be negatively impaired. In the last sections, we analyze the data center simulation results and evaluate them afterwards.

II. Related Work

The authors of [8] challenge the general assumption of an almost inelastic electricity demand. Based on innovations in information technology (IT) sector, electronics cost dropped. This would allow the integration into every day devices (e.g. Heating, Ventilation and Air Conditioning (HVAC), water heating, etc.) so that they can react to externally price or shed signals to manage the power consumption and maintain more stable market prices. Offering multiple approaches for their load control, their idea based on a central management authority that connects to every device. However, they recognized the need to adapt the power adjustment depending on the device's specific usage profile as the acceptance of load controlled devices depends on unaffected comfort of the users.

In [9], data center's power consumption was considered as regulation capacity comparable to an energy storage. Exploiting the constraints of the workload creates regulatory capacity. Utilizing an external regulation signal and dynamic voltage and frequency selection of the processors allows the adjustment of power consumption within set boundaries. Without violating quality of service the authors assumed an average regulation capacity of 8% of the data center's over all power consumption, depending on the server utilization. Compared to our approach, they only applied the signal periodically, dependent on an externally provided power budget signal which determines the servers frequency.

Comparable to our work the authors of [3] try to ease the integration of renewable energies into the power grid by handling data centers as large, flexible loads with a shiftable workload. They research various types of demand response approaches and compare it with energy storage like [9]. Duties, risks, and benefits of demand response were analyzed to leverage hidden opportunities for data center operators and and energy providers for additional value of the data center. Compared to this approach, we take advantage of data center capabilities as an addition to the secondary function of power plants, the provision of positive and negative primary control.

In [10], multiple degrees of freedom for the regulation of power consumption were researched. Their approach minimizes the power consumption of the data center while maintaining the guaranteed Quality of Service (QoS). Using their evaluated design of an architectural framework, they optimized the tradeoff concerning the number of virtual machines, their allocated central processing unit (CPU) cores, and the CPU frequency scaling for energy efficient cloud computing. Their results show that they still fulfill Service Level Agreements (SLA) while reducing the power consumption. The energy efficiency approach must be distinguished from our claim to grid stabilization. While energy efficiency reduces the overall power consumption, demand peaks or power plant outages may still happen which can be covered by our decentral load control idea as we reduce the power consumption autonomously based on the current grid state which would still be uncovered by [10].

The utilization of P-States for power control respectively power capping was researched in [11]. Compared to the switching of servers, DVFS excels with minimum overhead and fast response. They designed a two stage controller to control the P-States of a server with regard to the power consumption in one loop and a performance loop to ensure the SLAs. Despite their interesting results, their system relies on a control instance that manages the overall power consumption by implementing monitoring tools into the virtual machines. They target at power consumption and cost minimization without taking the power grid state into account. Our approach utilizes available resources to apply only minimal modification, so that we interact with the available frequency selection and scaling rather than calculating frequencies on our own. One advantage is a significantly faster frequency response time. As shown in [10] DVFS promises overall positive effects in power consumption adaptation of servers even though it is necessary to take SLAs and QoS into account.

III. Decentral Load Control and Demand Response

Data centers fulfill the necessary condition for power consumption regulation if at least parts of their workload can be shifted. This condition is important as it is rarely possible to drop workload as data center operators are bound to provide Quality of Service and fulfill Service Level Agreements. Various studies investigated how data centers could temporarily reduce their power consumption. Possible variants include the reduction of the number of servers used ([10], [12], [13]), the adjustment of the temperature set point for the climate equipment ([14]), or the utilization of the power saving features of the processors ([9]–[11], [15], [16]).

The authors of [3] mention various challenges for demand response, although some aspects do not apply to our approach of autonomous power regulation. At first, necessary investments for power consumption regulation. Decentral load control costs are considerably low as only a local frequency detection within a data center needs to be installed. A sufficient frequency measurement can be realized by using a reference clock count. Associated circuits, description and evaluation can be found in [17]. Servers obtain the measured grid frequency over internal network structure for further processing within an additional control software. This allows a fine-grained control approach comparable to the primary control of conventional power plants. Thus, not requiring an external control signal from the utility provider, and still retaining the authority over the data center's operation. Second, demand response issues coordination problems, as different electricity providers might decide differently the way data centers should adapt their power consumption. Our approach relies on a local grid frequency measurement as adaptation signal. The grid frequency is the same throughout the entire grid except for a time offset and small local variances. This allows a similar behaviour of grid supporting data centres in the power grid. Beneficial of this method is that it avoids the liability question for imbalances of faulty adaptation. Another challenge is the participation of data center operators. Either they participate voluntarily due to the advantages mentioned, on the basis of incentives, such as adapted electricity rates or other financial bonuses, or they could be encouraged via the already mentioned *Code of Conduct for Data Centers*. For *demand response* financial aspects often dominate as electricity rates increase with high power usage [4], [18]. The cost increase forms an interest to participate in demand response to reduce the operational cost of the data center. Even grid operators benefit from demand response as their operational cost decrease. Still, our intention only focuses on the power grid stability which does not necessarily correlate with electricity rates. Comparable to the obligation of power factor correction for electronic consumers with more than 75W power loss to reduce reactive power, there could be regulatory requirements to improve grid stability.

IV. DVFS FOR GRID STABILIZATION

The efficiency of data centers is often compared via the Power Usage Effectiveness (PUE) value. This indicates how much power is used for the operation of IT systems compared to the total consumption. Ideally, the PUE is 1.00. In reality, it is between approx. 2 and 1.1 depending on the size of the data centers, whereby larger data centers are generally considered more efficient and therefore have a value closer to 1 [19]. As the majority of power consumption in the IT systems arises from the processors of the servers we primarily target them. Secondary, this also affects other elements of a data center such as cooling or energy supply.

To minimize interference with the data center operation we utilize the *Dynamic Voltage and Frequency Selection* (DVFS) functionality. A performance state is selected from a processor-dependent list which determines both the necessary operating voltage and the clock frequency. The P-state with the highest performance, P0, uses the highest voltage for the highest possible frequency. With increasing number, P1 - Pn, voltage and frequency decrease to reduce the power requirement at lower performance. This selection is either done by the operating system or, in modern processors, in hardware (Hardware P-States). Based on desired behaviour, different methods can be used to more ore less aggressively select the necessary P-state, named for example *power saving* or *performance*. All cores on the processor share the same voltage, therefore also the P-state. The highest necessary P-state selected is used independent of the load of the remaining cores. A high P-state without load only marginally increases the power consumption. Modern processors, under certain circumstances, boost their frequency above the base frequency. The base frequency designates the clock rate which can be maintained by all cores under full load. The dynamic boost approach allows the redistribution of resources for higher frequencies on fewer cores below full load. This depends like power budget, thermal capacity, processor binning, silicon quality, and utilization.

Based on a local power grid frequency measurement, we manipulate the P-state selection by limiting the range of states a processor can select from. A correct adjustment of the P-state provides primary control power by withholding a small part of the computing power. The reservation of some of the highest states even allows the provision of negative primary control energy when used as soon as the grid state requires it. However, we will not address the issue of responsibility that arises from following our grid regulation approach as this is out of scope of our work and up to practical implementation details and regulatory decisions.

A. Test setup and prototype

Dependent on the current power grid frequency, our method adjusts the upper CPU frequency selection (in compliance with the supply voltage) as shown in Figure 1. Hence, the integration of a butterworth lowpass filter of first order and a cutoff frequency of 0.125 Hz effectively reduces grid frequency noise spikes and smoothes the power consumption adaptation, only adding a small, neglectable delay. Control oscillations, could happen as the proposed control only uses the power grid frequency which typically oscillates with frequencies between $0.2Hz$ - $1.5Hz$ [20]. But these frequency oscillations get damped by by Power System Stabilizers within power plants. The CPU speed we use in Figure 1 and in the following describes the maximum allowed CPU frequency in relation to the hardware limited maximum CPU frequency. While the power grid frequency is within a $\pm 20mHz$ around the norm frequency of the grid, the CPU frequency is set to a maximum performance of 90%. Under the condition that the grid frequency deviation increases above the $+20mHz$, the upper CPU frequency limit increases and provides negative control power. If the grid frequency falls below $-20mHz$, the upper processor limit reduces linearly to achieve a corresponding adaptation to the grid. At a grid frequency deviation of $-75mHz$ and further dropping frequency the negative slope increases and reduces the upper CPU limit faster. At a mains frequency deviation of $\pm 100mHz$ our adaptation holds a constant value of 70% or 100%, respectively. Usually, the grid stabilization requires positive control power, so that we designed this algorithm asymmetrical, to be able to provide more positive control power than negative.

Often, power consumption simulations of processors use a linear model, similar to Eq. 1. Due to the boost feature this model does not apply anymore. Therefore, we measured the power consumption of a state-of-the-art processor by stepwise increasing the number of fully loaded cores by using the linux tool *stress*. Our test computer uses an Intel Core i7 9700 CPU with 8 cores and 8 threads. The power consumption measurement is realized by a GW-Instek GPM8213 power meter. Figure 2 displays the result of this measurement, normalized to the maximum power consumption, and confirms our statement that the power requirement is not linear.

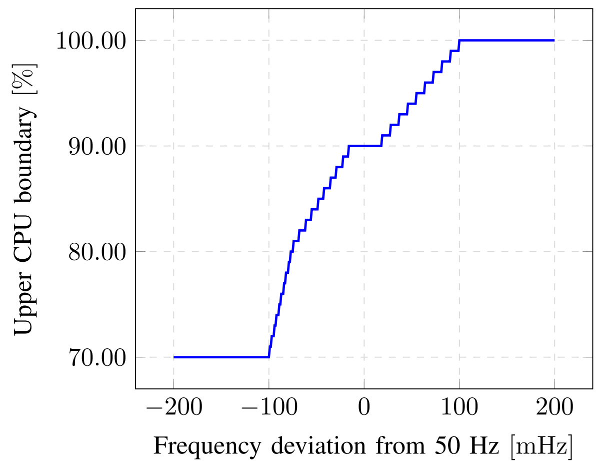

Fig. 1. Upper CPU frequency boundary dependent on the main frequency deviation

$$P = P_{idle} + (P_{maximum} - P_{idle}) * u \quad (1)$$

As written before, we want to manipulate the limits of the P-State selection. Since we do this in a range of $70\% - 100\%$ of the clock frequency as shown in Figure 1, we performed further measurements where we set the measured grid frequency deviation to fixed values between $-100mHz$ to $100mHz$ in $20mHz$ steps. This test uses the same workload scenario, gradually increasing core utilization. Figure 3 shows the relative power consumption normalized to the maximum power consumption, depending on the aforementioned fixed grid frequency values and number of cores under full load. We observe a flattening of the curve, resulting in a relationship between processor utilization and power demand that approaches a linear function at around 80% of the maximum frequency. Reducing the limit even further bounds the maximum power consumption under full load. We expect this to be the general behaviour of modern processors because most current processors have boost features. Naturally the specific function between power consumption and utilization may vary dependent on the processor.

We suspect that despite positive results, our presented algorithm will find only limited acceptance

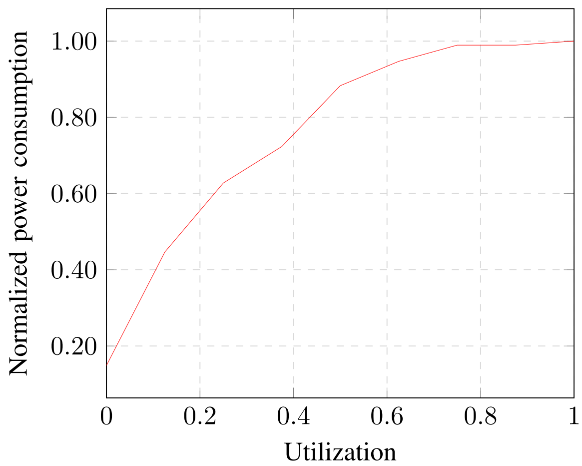

Fig. 2. Normalized power consumption of an Intel Core i7 9700 dependent on the utilization

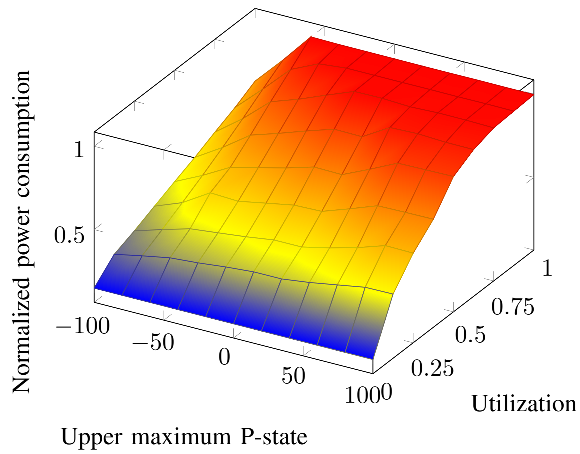

Fig. 3. Normalized power consumption of an Intel Core i7 9700 for different loads and varying allowed maximum P-state

due to the limitation of the computing capacity to 90% for the target grid frequency of $50Hz$. Reasons for this would be, for example, that computing power is reserved for a few times a year and it depends on the workload whether negative control power is fed in.

For practical reasons, it seems more realistic that data center operators adopt a regulation that only provides positive balancing power. In this case, the maximum computing power would be available at and above $50Hz$ grid frequency, but the P-state list and thus the power consumption would be reduced when the grid frequency drops. The characteristic curve from Fig. 1 could be shifted so that the highest P-state is available at 49.98 Hz.

From an academic point of view, however, it still seems worth investigating due to the characteristics of data centers, since only few loads are capable of providing negative control power without necessarily influencing operation.

In the next section we present our data center simulations, in which we have integrated the above mentioned data for more realistic results by linearly interpolating missing data.

B. Data center for load control

To analyze our power consumption adaptation, we used MATLAB Simulink to build the data center power consumption model according to the description and parameterization specified in [21]. They modelled the major factors of power consumption of a data center:

- *Processors*
  The data center's processor model in [21] used Equation 1 for the power consumption calculation. We have integrated our measurements by replacing the utilization factor u in Equation 1 with a linear interpolation of the measurement results using the current processor utilization and maximum processor speed from Figure 3.
- *HVAC*
  The power consumed by the processor converts to heat which must be transported away from the servers. Typically, data center operators use a Computer Room Air Conditioning (CRAC) or Computer Room Air Handler (CRAH) system to cool their server farm. While CRAC units work comparable to an air conditioner by blowing air over cooling coils filled with refrigerant, CRAH units, instead of refrigerant, use chilled water provided by chillers. Using CRAH units usually consumes less power compared to CRAC as the air cooled by the environment can be used and therefore increases the data center's efficiency, also called Power Usage Efficiency (PUE).

The following simulations use CRAH cooling due to its higher efficiency.

- *Power supply and power distribution*
  The Uninterruptable Power Supply (UPS) allows the bridging of power failures. Positioned between the main grid and the data center to seamlessly maintains the supply of the data centre in an emergency. The Power Distribution Units (PDU) connect every server to the power supply of the UPS. Both, power supply and distribution to the server farm may involve significant load-dependent losses.

For further details of the data center model, we refer to [21].

An important condition for the reliability of our simulation results consists in the ability to shift workload. Furthermore, we do not consider additional costs arising from workload shifting, e.g. moving workload across data center resources. Depending on the grid situation our algorithm adapts the P-state selection range. This directly affects the maximum utilization rate of the servers. Typical mean utilization rates for data centers are far below 100%, for example hyper-scale data centers have the highest mean utilization rate of 50% [19].

In [21], an utilization rate of a time period of one week is reported which we reuse within our simulations. This utilization of the data center represents the workload that needs to be processed. Reducing the maximum frequency of the processor does not necessarily shift workload because at first only the effective utilization rate rises. If the current workload increases beyond the current processing capacity, workload gets postponed until processing capacities become available.

Equation 2 to 4 show the principles of the calculation of the workload to be processed or the utilization, respectively. Basically, the workload that needs to be processed w_{total} consists of the regular and the shifted workload. As the algorithm adjusts the P-state selection, the maximum processable workload ($w_{processable}$) is capped affecting the maximum utilization rate. The data center might need to postpone some workload ($w_{shifted}$) to the next time slot when processing capabilities become available.

$$w_{total}(t) = w(t) + w_{shifted}(t) \tag{2}$$

$$w_{shifted}(t+1) = w_{complete}(t) - w_{processable}(t) \tag{3}$$

$$w_{processable}(t) = w_{total}(t) * d_{adaptation}(t) \tag{4}$$

The parameter $w_{complete}$ represents the overall processed workload that, integrated over time, allows the comparison with a reference data center without *decentral load control*. At the end of the simulation the integrated $w_{absolute}$ should be equal. Otherwise, the power adaptation of decentral load control uses an suboptimal configuration with too aggressive processor capping.

The calculation of the power consumption of the data center's processors uses $w_{complete}$ in combination with the processors power consumption behaviour of Figure 3. For a proper modelling, we exchanged the factor *u* of Equation 1 with the results of Figure 3 that depend on utilization and grid frequency.

V. RESULTS

We initially tested the model with a sine-wave signal as artificial power grid frequency, viewable on the right axis of Figure 4. This allows us to to investigate the effect of larger frequency deviations which only happen a few times a year. The power consumption of the decentral load control enabled data center can be compared to a reference data center without decentral load control in Figure 4. During the negative sine phases, the processors capacity gets reduced, to limit the data center's power consumption. Due to this reduction, also less power is necessary for cooling and power distribution increasing the beneficial effect of using decentral load control. In time interval 15 - 20 the data center's overall power consumption reduces by up to 11%. During the negative sine phases, the processors capacity gets reduced, to limit the data center's power consumption. Due to this reduction, also less power is necessary for cooling and power distribution increasing the beneficial effect of using decentral load control. In time interval 15 - 20 the data center's overall power consumption reduces up to 11% by shifting workload to a later time. Without workload shifting during the negative sine phases, the processors maximum P-state reduces without impacting the processed workload but a maximum power consumption reduction of 7.6% due to more efficient P-states, observable in time interval 64 - 70. During the positive sine wave phases of the artificial grid frequency our adaptation enables higher P-states. Despite the usual workload it also also allows to process shifted workload or reduces the necessity for workload shifting, respectively. In time interval 22 - 23 of Figure 4 shifted workload processing increases the power consumption up to around 19.9%. Everytime the power consumption increases compared to the reference data center relates to an increase of the absolute workload due to workload shifting, see Figure 5. As the overall processed workload equals to the reference data center, we can safely assume that we will achieve a positive power grid stabilization effect.

Larger frequency variations, e.g. larger than $200mHz$, usually happen only few times in the year. Common are small frequency deviations up to 100 mHz. Figure 6 and 7 present the power consumption and workload processing based on a grid frequency data measurement without any grid incidents, shown on the right axis of the former Figure. The observable noise results from different time resolutions of the grid frequency measurement (resolution: 1s) and the utilization rate (resolution: 1h).

Using our decentral load control method slightly reduces the overall power consumption of the data center by shifting the workload to more efficient P-states of the processors. In Figure 6, around the time 20 and 68 small workload shifting shows this for an

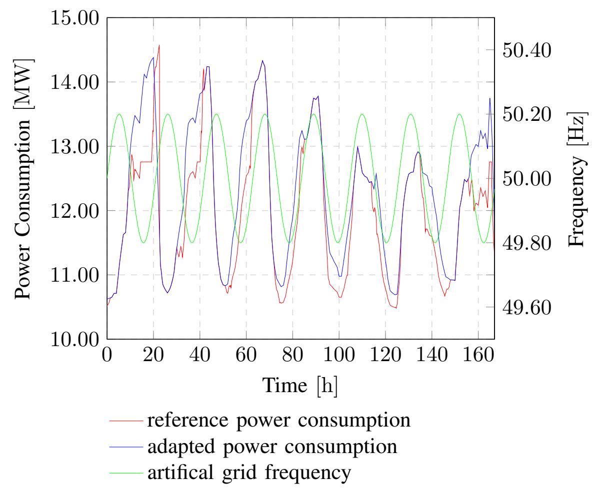

Fig. 4. Power consumption comparison between the reference data center model and main frequency adapting data center with an artificial sine wave frequency

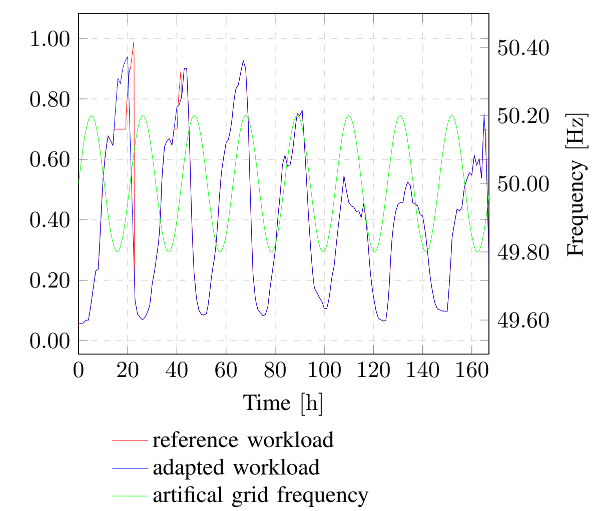

Fig. 5. Workload comparison between the reference data center model and main frequency adapting data center with an artificial sine wave frequency

everyday scenario. Without affecting the workload processing a maximum of 7.7% of the overall power consumption could be reduced. Therefore, data center operators with utilization rates below their current processing capacities benefit from reduced frequencies. The highest power consumption difference for our example grid frequency data are around 9.4% compared to the reference data center, while it was also possible to increase the power

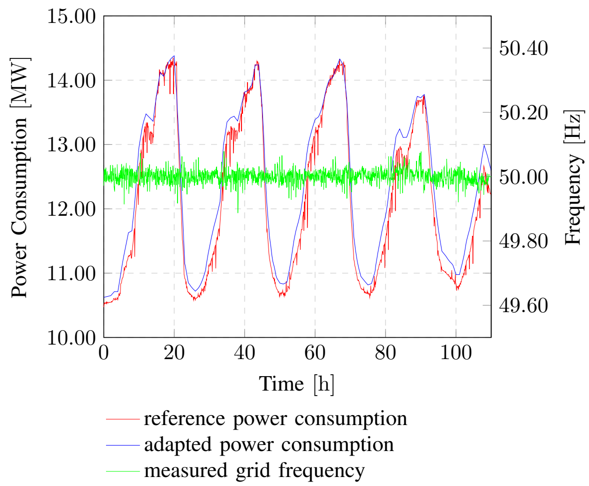

Fig. 6. Power consumption comparison between the reference data center model and main frequency adapting data center with a measured main frequency data

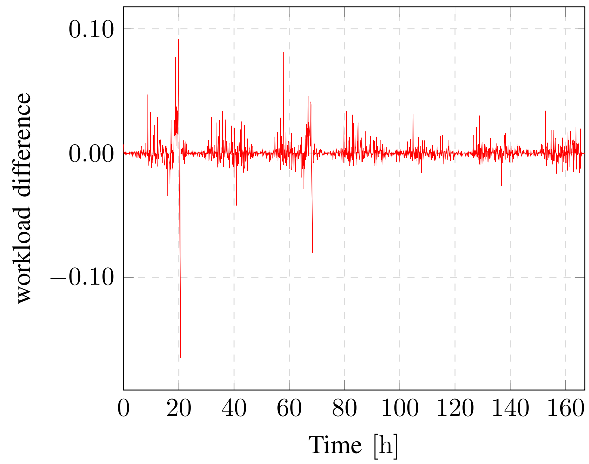

Fig. 8. Workload difference between reference model and main frequency adapting model

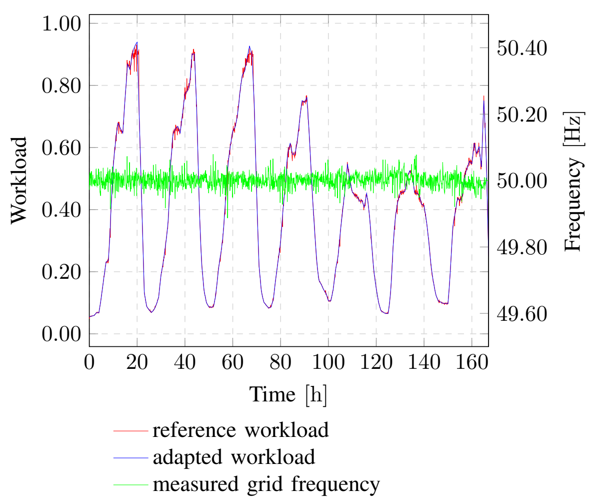

Fig. 7. Workload comparison between the reference data center model and main frequency adapting data center with a measured main frequency data

consumption by up to 3.8% due to workload shifting. A potential beneficial effect of the latter one depends completely of the current grid situation.

VI. Conclusion

With exceptionally low costs, based on a single mains frequency measuring unit and a server-side software, the operation of data centers can be extended by a power grid oriented function. Not only do they provide positive but also negative control power at least as far as some conditions, e.g. shiftable workload, apply. Due to the stable high share of the total energy demand, a new type of cost-effective primary control energy could be established in the future. However, this requires the data center operators to participate in the grid stabilization process. In this context, remuneration for the provision of balancing power should be provided, similar to that for power plants, so that the operators participate voluntarily. Several incentive systems, already familiar by demand response, are possible, such as tax reductions or bonuses from the grid operators. The future grid, in which conventional power plants will gradually be shut down, requires sufficient control power for safe operation, so it seems sensible to exhaust all realistic stabilization possibilities. Due to the increasing decentralization of power generation due to wind and water, our approach also leads to a decentralized provision of necessary system services to ensure a stable power grid through a broad geographical distribution and at the same time synchronicity due to the use of the local grid frequency. Our work confirms the positive effect on the grid stability and demonstrated that the PUE can be increased even during normal operation and that possible incentive systems can also be used. All Matlab Simulink models and simulations, complete with parameters can be found at [22].

ACKNOWLEDGMENT

This paper was made within the framework of the research project "Netz-Stabil" and financed by the European Social Fund (ESF/14-BM-A55-0023/16). This paper is part of the qualification program "Promotion of Young Scientists in Excellent Research Associations - Excellence Research Program of the State of Mecklenburg-Western Pomerania".

REFERENCES

[1] Department of Energy, "The smart grid: an introduction," 2008, prepared for the U.S. Department of Energy by Litos Strategic Communication under contract No. DE-AC26-04NT41817, Subtask 560.01.04,.

[2] C. Greer *et al.*, "NIST framework and roadmap for smart grid interoperability standards, release 3.0," National Institute of Standards and Technology, Tech. Rep., oct 2014.

[3] A. Wierman, Z. Liu, I. Liu, and H. Mohsenian-Rad, "Opportunities and challenges for data center demand response," in *International Green Computing Conference*. IEEE, nov 2014.

[4] Z. Liu, A. Wierman, Y. Chen, B. Razon, and N. Chen, "Data center demand response: Avoiding the coincident peak via workload shifting and local generation," *Performance Evaluation*, vol. 70, no. 10, pp. 770–791, oct 2013.

[5] F. Uster and D. Timmermann, "Decentral load control for grid stabilization," in *2019 IEEE 17th International Conference on Industrial Informatics (INDIN)*. IEEE, jul 2019.

[6] M. Avgerinou, P. Bertoldi, and L. Castellazzi, "Trends in data centre energy consumption under the european code of conduct for data centre energy efficiency," *Energies*, vol. 10, no. 10, p. 1470, 2017.

[7] A. Ulbig, T. S. Borsche, and G. Andersson, "Impact of low rotational inertia on power system stability and operation," *IFAC Proceedings Volumes*, vol. 47, no. 3, pp. 7290–7297, 2014.

[8] M. Ilic, J. Black, and J. Watz, "Potential benefits of implementing load control," in *2002 IEEE Power Engineering Society Winter Meeting. Conference Proceedings (Cat. No.02CH37309)*. IEEE, 2002.

[9] S. Li, M. Brocanelli, W. Zhang, and X. Wang, "Data center power control for frequency regulation," in *2013 IEEE Power & Energy Society General Meeting*. IEEE, 2013.

[10] S. Tesfatsion, E. Wadbro, and J. Tordsson, "A combined frequency scaling and application elasticity approach for energy-efficient cloud computing," *Sustainable Computing: Informatics and Systems*, vol. 4, no. 4, pp. 205–214, dec 2014.

[11] X. Wang and Y. Wang, "Coordinating power control and performance management for virtualized server clusters," *IEEE Transactions on Parallel and Distributed Systems*, vol. 22, no. 2, pp. 245–259, feb 2011.

[12] D. J. Bradley, R. E. Harper, and S. W. Hunter, "Workload-based power management for parallel computer systems," *IBM Journal of Research and Development*, vol. 47, no. 5.6, pp. 703–718, sep 2003.

[13] O. Sarood, A. Langer, L. Kale, B. Rountree, and B. de Supinski, "Optimizing power allocation to CPU and memory subsystems in overprovisioned HPC systems," in *2013 IEEE International Conference on Cluster Computing (CLUSTER)*. IEEE, sep 2013.

[14] C.-J. Tang, M.-R. Dai, C.-C. Chuang, Y.-S. Chiu, and W. Lin, "A load control method for small data centers participating in demand response programs," *Future Generation Computer Systems*, vol. 32, pp. 232–245, mar 2014.

[15] T. Komoda, S. Hayashi, T. Nakada, S. Miwa, and H. Nakamura, "Power capping of CPU-GPU heterogeneous systems through coordinating DVFS and task mapping," in *2013 IEEE 31st International Conference on Computer Design (ICCD)*. IEEE, oct 2013.

[16] T. Kolpe, A. Zhai, and S. S. Sapatnekar, "Enabling improved power management in multicore processors through clustered DVFS," in *2011 Design, Automation & Test in Europe*. IEEE, mar 2011.

[17] F. Plocksties, "Stable Grid," 2017, GitHub repository. [Online]. Available: https://github.com/franzpl/StableGrid

[18] I. Koutsopoulos and L. Tassiulas, "Challenges in demand load control for the smart grid," *IEEE Network*, vol. 25, no. 5, pp. 16–21, sep 2011.

[19] A. Shehabi, S. Smith, D. Sartor, R. Brown, M. Herrlin, J. Koomey, E. Masanet, N. Horner, I. Azevedo, and W. Lintner, "United states data center energy usage report," 6 2016.

[20] Holger Berndt, Mike Hermann, Horst D. Kreye, Rüdiger Reinisch, Ulrich Scherer, Joachim Vanzetta, "Transmission Code 2007 / Netz- und Systemregeln der deutschen Übertragungsnetzbetreiber," Verband der Netzbetreiber (VDN), Tech. Rep., 2007.

[21] R. Rahmani, I. Moser, and M. Seyedmahmoudian, "A complete model for modular simulation of data centre power load," *arXiv preprint arXiv:1804.00703*, 2018.

[22] F. Uster, "Netz-Stabil - Dezentrale Netzstabilisierung," 2020, GitLab repository. [Online]. Available: https://gitlab.com/FelixUster/netz-stabil---dezentrale-netzstabilisierung