[·] Robert Basmadjian, Florian Niedermeier, Gergoe Lovasz, Hermann De Meer and Sonja Klingert. GreenSDAs leveraging power adaption collaboration between energy provider and data centres. In *Proc. of the 3rd IFIP Conf. on Sustainable Internet and ICT for Sustainability (SustainIT 2013)*. IFIP, Oct. 2013.

See `http://www.net.fim.uni-passau.de/papers/Basmadjian2013b` for full reference details (BibTeX, XML).

GreenSDAs Leveraging Power Adaption Collaboration between Energy Provider and Data Centres

Robert Basmadjian
University of Passau
Innstrasse 43, Passau, Germany
Email: basmadji@fim.uni-passau.de

Florian Niedermeier
University of Passau
Innstrasse 43, Passau, Germany
Email: niederme@fim.uni-passau.de

Gergő Lovasz
University of Passau
Innstrasse 43, Passau, Germany
Email: lovasz@fim.uni-passau.de

Hermann De Meer
University of Passau
Innstrasse 43, Passau, Germany
Email: demeer@fim.uni-passau.de

Sonja Klingert
University of Mannheim
A5, B101, Mannheim, Germany
Email: klingert@informatik.uni-mannheim.de

***Abstract*—Data centres, due to their significant energy use and highly automated IT infrastructure, are excellent candidates to participate in demand response programs. However, the major inconvenience of today's electricity tariffs between energy provider and its customers is their lack of flexibility, which renders demand response programs difficult to realise. In this paper, we propose a new supply demand agreement in order to foster power adaption (i.e. increase/decrease) collaboration between energy provider and data centres. To this end, we introduce contractual terms and based on those we propose reward and penalty schemes. Furthermore, we provide a signalling scheme which defines the communication requirements necessary to enable power adaption collaboration. Finally, we present a scheduling policy which helps the energy provider to request data centres in a fair manner for power adaption.**

I. INTRODUCTION

Demand response (DR) is a mechanism used in electricity grids to manage customers' energy consumption during power shortage (i.e. peak) situations. It consists of a list of actions that need to be taken by the customers in order to reduce the electrical load when the power grid is congested. Such a mechanism is crucial in order to, on one hand ensure the stability of electricity grids, and on the other guarantee no sudden change takes place in the market conditions causing the prices of electricity cost to increase.

In [5], the Demand Response Research Center (DRRC) at Lawrence Berkeley National Laboratory analysed data centres (DCs) for DR. The results show that DCs, on the basis of their operational characteristics and energy use, have significant potential for DR. Furthermore, they demonstrated that the key efficiency technologies for DCs to participate in DR are strategies to reduce cooling energy use as well as virtualisation. In order to assess the results of [5], the authors in [3] and [4] show that through energy-aware optimisation policies, it is possible to reduce between 17% – 20% the energy use of a private cloud.

As mentioned above, numerous incentives exist for DCs to participate in DR mechanisms. However, the major inconvenience in today's market condition is that the electricity tariffs between energy provider (EP) and its customers do not allow for DR mechanisms to take place due to their inflexibility. Another aspect which is becoming more and more prominent in our daily life are renewable energy sources (e.g. photovoltaic). The major disadvantage of such renewables is their intermittent behaviour which requires careful and quick planning by the EP in order to ensure the stability of the electricity grid. Hence, today there is a necessity for new supply demand agreements to be incorporated into the market that consider the customers' flexibility in both reduction/increase of their power consumption.

In this paper, we propose a new supply demand agreement that leverages power adaption collaboration between an EP and its customers (e.g. DCs). We call such an agreement *GreenSDA* which is introduced in Section IV. More precisely, we specify the agreement terms needed by DCs to participate in DR mechanisms and based on those terms we propose *reward* and *penalty* schemes. In order to enable DR mechanism between an EP and DCs, in Section V we introduce the signalling schemes by taking into account our GreenSDA. In Section VI, we present a *DC scheduling policy*. The main objective of such a policy is to select appropriately DCs whenever power surplus/shortage situations take place in the electricity grid. In Section VII, we provide a *discrete-event simulation* which analyses a one year grid state of a major German EP, E.ON[1]. The simulation evaluates the effectiveness of the overall approach as well as it demonstrates the quality of load balancing of the proposed DC scheduling policy.

II. RELATED WORK

One important topic in connection with DR approaches is how the incentive is structured. In general, the incentive

---

[1] http://www.eon.com/en.html

discussed in theory and applied in practice is based on reduction of the electricity bill. There are two fundamentally different approaches to this respect: on one hand, incentives for DR can be seen as topics connected with the pricing of electricity (e.g. price-based programs or Time of Use (ToU) pricing); on the other hand, these incentives can be granted in the form of a reward that is deducted en bloc from the electricity bill (e.g. incentive-based program).

An overview of these concepts is given in [1] and [2]: ToU electricity pricing is an established instrument for subtle steering of electricity demand. Other variants of price-based programs according to [1] are Critical Peak Pricing (CPP) and Real Time Pricing (RTP). RTP is the best option as it is the market or cost-based price that leads to a market equilibrium that perfectly levels demand and supply. However, it is also the option that is the most complex to carry out. Incentive-based programs include among others Interruptible/Curtailable programs (I/C), Direct Load Control (DLC), and Emergency Demand Response Programs (EDRP). At I/C programs, customers qualifying through a certain minimum electricity demand (e.g. 200kW or 1MW) get a bill credit if they reduce load in critical situations – and are penalised if they don't respond. DLC gives a lot of control to the EP who at the expense of an energy bill reduction can remotely shut down specific sites at the customer (e.g. heating equipment). All these approaches however, do not foster a dynamic adaption process at the side of the customer like the one suggested in this paper. The presented approach allows for a much more fine-granular power adaption that smoothly matches electricity demand to supply instead of crudely shedding more load than necessary. Also, unlike the proposed GreenSDAs of this paper, the adaption to intermittent renewable energy sources has not yet been a topic of DR incentives.

From a technical point of view, there are some advanced standards for DR management in the electricity grid. These are, e.g. OpenADR Communication Specification (version 1.0) [9] and OASIS Energy Interoperation (version 1.0) [8]. The OpenADR standard is an open DR communication standard for data exchange between consumers, EPs and independent system operators. OpenADR defines a system architecture for DR management, data models for DR-related information and communication interfaces between the participating entities. After subscribing to a DR program of an EP, the customer registers resources (e.g. lighting, air conditioning) at the Demand Response Automation Server (DRAS) that plays the role of a negotiator between the customer and the EP. For each resource, the customer can define a set of constraints and schedules. The constraints describe capabilities of the resource, like the amount and the duration of the maximum power adaption the resource is able to perform. The schedules contain rules that define the reaction of the resource to adaption requests (called DR Event) of the EP. When the EP sends a DR Event to a customer, it has to consider the constraints of the registered resources. A DR Event usually contains time intervals with corresponding (resource specific) pricing information but may also contain direct load control instructions. The reaction of a resource to different pricing levels is defined in the schedules of the resource. Even though OpenADR represents sophisticated DR approaches, it does not explicitly address DCs. However, DCs have special requirements regarding the integration into DR approaches. The power adaption capability of a DC, e.g., might be highly dynamic due to changing workload patterns. Therefore, the EP must be informed on the changing adaption potential of the DC. Furthermore, methods applied for decreasing power consumption (e.g. using the UPS instead of the power grid) may entail a significant increase of power consumption after the adaption when the DC recovers and returns to normal operation (e.g. the UPS has to be recharged). Therefore, information on power consumption during recovery has to be considered when DCs are integrated in DR. Whether and how OpenADR supports DR with DCs is unclear. The signalling scheme suggested in this paper addresses explicitly the integration of DCs into DR and considers DC-specific requirements.

To the best of our knowledge, this is the first approach that targets Demand Response and power adaption collaboration between energy provider and data centres. Hence, the proposed GreenSDA is novel to the literature.

III. PRELIMINARIES

In this section, we provide definitions and assumptions considered throughout the rest of this paper.

Definition 1: We define a power surplus as an abundance of available power, which can be expressed as a deviation of the ratio $\frac{powergeneration}{powerdemand}$ towards values > 1.

Definition 2: We define a power shortage as a scarcity of available power, which can be expressed as a deviation of the ratio $\frac{powergeneration}{powerdemand}$ towards values < 1.

Assumption 1: We assume that the EP has means to detect potential power surplus/shortages.

Assumption 2: We assume that neither EP nor the DCs are controlled by malicious operators.

Assumption 3: We assume that each DC has means to increase/decrease its power consumption. Among those, the following are the most relevant ones whose details are out of the scope of this paper:

1) Workload consolidation and job shifting to reduce power consumption.
2) Heating up (cooling down) the DC in order to reduce (increase) the power consumption.
3) Run the DC on battery (e.g. UPS) to decrease the consumption and use the battery of the UPS as an energy storage to increase the power consumption.

Assumption 4: We assume that every DC that participates in DR has a unique ID. The communication between the EP and DCs is assumed to be reliable and secure.

Assumption 5: We assume that every DC participating in DR is equipped with a smart meter. It is required that at least the EP has access to the smart meter. DC's access to the monitoring data of the smart meter is optional.

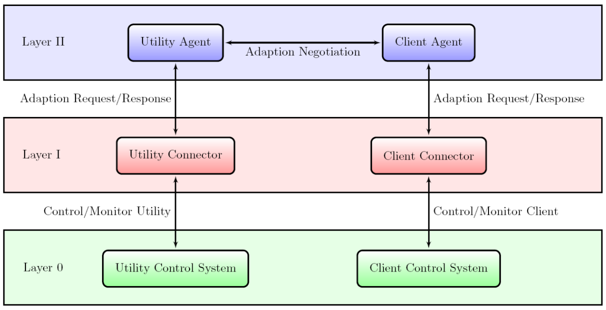

Fig. 1. Demand Response architecture

Assumption 6: Throughout the paper, we assume that the parties involved in the DR system are an EP and several DCs. However, the developed system is flexible enough to handle other utility/client combinations. Fig. 1 shows the generic architecture of the developed utility/client interaction pattern. To adapt the system to different utilities or clients, only the connector components located at Layer I have to be modified to match the implementation of control and monitoring systems at Layer 0. The algorithms executed on Layer II are agnostic of this adaption.

IV. Green Supply Demand Agreements

In this section, the proposed Green Supply Demand Agreement (GreenSDA) between the EP and DCs is introduced. Then, based on the agreement terms, we give our reward and penalty schemes.

A. Contractual Terms

In order to foster collaboration between the EP and DCs for power adaption purposes, the current electricity tariffs need to be altered. The main reason for introducing new kinds of energy contracts is that the current tariffs lack flexibility in terms of power adaption. To this end, our proposed GreenSDA includes contractual terms allowing for power adaption collaboration between an EP and DCs. Note that the proposed terms are generic enough that can be applied to any type of customers (e.g. household, industry):

Term 1: **minIncrease**

1) minDurationMinIncrease
2) maxDurationMinIncrease

Term 2: **maxIncrease**

1) minDurationMaxIncrease
2) maxDurationMaxIncrease

Term 3: **minDecrease**

1) minDurationMinDecrease
2) maxDurationMinDecrease

Term 4: **maxDecrease**

1) minDurationMaxDecrease
2) maxDurationMaxDecrease

Term 5: **requestPeriod**

Term 6: **maxAdaptionTime**

Term 7: **maxReactionTime**

Term 8: **maxRejectsPerMonth**

Term 9: **maxRejectsInSuccession**

Term 10: **maxRequestsPerMonth**

Terms 1 and 2 (Terms 3 and 4) represent the minimum and maximum power increase (decrease) in kW that the DC can adapt at any given point in time. For each of the above mentioned power adaption capabilities, the minimum and maximum duration expressed in minutes need to be specified. Term 5 denotes the period during which the EP can send an adaption request to the DC. Among possible values of this term are *morning/noon/evening/night/midnight*. Term 6 indicates the maximum time (minutes) the DC needs to adapt its power consumption. This term shows how quickly the DC adjusts its energy use. Once the EP sends a power adaption request (e.g. increase/decrease) to the DC, Term 7 guarantees that this latter replies back within a maximum time (minutes). Usually, the values of such a term must not exceed 5 minutes. Since it is possible for a DC to reject the power adaption request of an EP, Term 8 specifies the maximum number of requests the DC can reject per month. Furthermore, Term 9 defines the maximum number of successive requests the DC can reject. Finally, Term 10 indicates the maximum number of power adaption requests an EP can send to the corresponding DC per month. Note that Terms 8–10 are on a monthly basis since these can have different values for various periods of the year. Furthermore, these terms need to be monitored in order to keep track of their actual execution in DCs. Consequently, *currentRejectsPerMonth*, *currentRejectsInSuccession*, and *currentRequestsPerMonth* are integrated to the monitoring systems of DCs (not part of GreenSDAs) in order to monitor the corresponding terms' actual fulfilment.

B. Rewards and Penalties in GreenSDA

Electricity tariffs generally consist of three parts: basic fixed fees, a power charge that varies with the (highest) power required, and an energy charge that depends on the kWh purchased. Similarly, GreenSDA reward and penalty schemes can contain both fixed and variable parts. As mentioned in Section II, there are both time- and incentive-based reward approaches to DR. In economic theory, of course only real time pricing of the cost met by demand leads to a perfect equilibrium. However, in a scenario closer to reality, it makes sense to combine time- and incentive-based approaches in order to create an atmosphere of reliability. Additionally, penalties in case of unexpectedly low collaboration should be foreseen. This leads to the following approach:

1) Components of the GreenSDA Reward Scheme: There are two different groups of terms that influence the reward on different levels of detail:

- *Static terms that are agreed upon at the time of contract signature and will not be altered over time; they represent the level of commitment of the collaboration and should thus be rewarded with a fixed amount.*
- *Dynamic terms that depend on the power shortage/surplus situations of the EP. Note that the proposed scheme is not a DLC program, where the EP has the right to remotely turn off certain equipment at the DC; rather the outcome relies on negotiation. Therefore a reward component is suggested that subsumes the variable terms and is based on a metric representing the collaboration between EP and DC. In case the collaboration is much lower than foreseen in the contract, a penalty is imposed.*

In sum, the GreenSDA reward and penalty balance can be calculated from the following three components:

$$Bonus = Reward_{Fixed} + Reward_{Variable} - Penalty. \quad (1)$$

2) Reward and Penalty Schemes:

a) The Fixed Reward: As mentioned above, the static reward ($Reward_{Fixed}$) is based on the commitment of the DC to reduce/increase its power consumption upon request of the EP. There are two options of how to design this: The static reward can either be a continuous function of the degree of commitment, or several categories can be created that reflect the degree of commitment.

An example for a continuous function is given by:

$$Reward_{Fixed} = Basic_{Fee} * \frac{\sum_i \alpha_i * CT_i}{100}, \qquad (2)$$

where $Basic_{Fee}$ denotes the power charge of the regular energy tariff. The parameter α_i indicates the weighting factor of the contractually agreed terms CT_i as described in Section 4.1. The weighting factors can be tuned such that the fixed reward is between 15% and 30%. However, it seems to be more practical to create categories like a bronze, silver, and gold GreenSDA. This approach is chosen for the current research, even though it slightly limits the flexibility of optional term combinations in the GreenSDA.

The idea is to determine bronze, silver and gold values for each term which then can be combined to a fixed reward. Thus the reward function is not continuous, but it still differentiates the DC's commitment to a great degree:

$$Reward_{Fixed} = Basic_{Fee} * \frac{\sum_i Gold + \sum_i Sv + \sum_i Bronze}{100}. \qquad (3)$$

In the presented example, the index i is one of the categories specified in Table I. Through these categories reward points are assigned, adding up to a percentage as in Equation (2). This can be done in the following way:

- Each gold is worth 6 points,
- Each silver is worth 4.5 points,
- Each bronze is worth 3 points

Applied to the Table I, this means that the fixed reward will be between 21% and 42% of the contracted power. This is calculated by adding up the points according to the example in Table I: If all 7 items are contracted in the gold version, they are valued 6 points each so that all together the reward comprises 42 points, which in the current example amounts to 42 % of the basic fee. This means that the DC contractually agrees to positively respond to at least 15 requests per year. In these cases, it either takes measures to reduce its power by 66% for up to two hours or to increase by 50 kW for up to 4 hours at each event (see rows 2 to 5 of Table I). The DC promises to effectually take these measures within 5 minutes at day and night, except for 3 hours that it specifies in the contract (see rows 6 and 7 of Table I).

b) The Variable Reward: The dynamic part of the GreenSDA reward scheme depends on the realised collaboration. Therefore, a metric for collaboration must be defined in a way that it can be monitored throughout the enacted collaboration:

$$Coll_{DC} = \frac{MWh_{AdapDC}}{MWh_{ReqEP}} * \frac{T - currentRejectsPerMonth}{T}, \qquad (4)$$

with $Coll_{DC} \in [0, 1]$ (a value of 1 indicates the highest collaboration). MWh_{AdapDC} and MWh_{ReqEP} represent the adaption (i.e. increased or decreased) effort of the DC and the power adaption requested from the EP respectively (in MW, MWh_{AdapDC} = requested power adaption duration). T denotes the total number of requests sent by EP to DC, whereas $currentRejectsPerMonth$ is defined in Section IV-A. There are at least two different implementation options: One option will affect the price per MWh for the current month (as the reward is given per MWh), the second option will relate only to the MWh shifted.

Option 1: The reward reduces the energy charge P_{DC} for all MWh consumed in the current month: $Reward_{Variable} = \beta * Coll_{DC} * P_{BASE} * MWh_{consumed}$ where $MWh_{consumed}$ gives the consumed energy use expressed in MWh, P_{BASE} denotes the regular electricity price per MWh, $Coll_{DC}$ indicates the degree of DC's collaboration as defined in Equation (4) and $\beta \geq 0$ is the weighting coefficient that needs to be specified in the GreenSDA.

Option 2: The reward changes the price only for the shifted MWh: $Reward_{Variable} = \beta * Coll_{DC} * P_{BASE} * MWh_{Reduced} + \gamma * Coll_{DC} * P_{BASE} * MWh_{Increased}$ where $\beta * P_{BASE} * MWh_{Reduced}$ computes the reward paid for the MWh avoided; something similar is realised in DR programs in the U.S., where β can go up to more than 1 (targeted at emergency situations). On the other hand, $\gamma * P_{BASE} * MWh_{Increased}$ calculates the reward paid for MWh integrated due to erratic supply from renewable. In both cases, $\beta \geq 0$ and $\gamma \geq 0$ are the weighting coefficients which are constant for a specific period.

c) Penalties: The penalty scheme relates to the fixed components as well as to the level of collaboration ($Coll_{DC}$) that the DC committed to. Each condition selected for the calculation of the basic GreenSDA has a value depending on the weighting coefficient. If any of these terms at the end of the billing period has been breached, the original reward on the basic fee has to be paid back. Additionally, a penalty depending on the realised collaboration $Coll_{DC}$ must be paid:

$$Penalty = \sum_i \rho i * CT_i + \delta(1 - Coll_{DC}),$$

with $Coll_{DC}$ indicates the degree of DCs' collaboration defined in such that $CollDC \leq a$ (with a being a threshold value specified in the GreenSDA), whereas $\alpha_i * CT_i$ is the originally granted reward on basic fee such that for the case of penalty $\rho_i \geq \alpha_i$. The parameter $\delta \geq 0$ denotes the weighting coefficient.

V. SIGNALLING SCHEMES

In order to realise the envisioned integration of DCs into a DR program, certain information needs to be exchanged between the EP and the DC. This section outlines a possible signalling scheme that can be used to integrate a DC – with highly dynamic power adaption capabilities, a complex site infrastructure consisting of several IT resources and significant recovery power consumption – into a DR program.

TABLE I. GOLD, SILVER, AND BRONZE GREENSDA CATEGORIES

Contractual Term (CT)	Gold	Silver	Bronze
maxRequestsPerMonth	if the sum of this term over the 12 months is ≥ 15	if the sum of this term over the 12 months is ≥ 12	if the sum of this term over the 12 months is ≥ 9
maxDecrease/contractedPower	$> 66\%$	33%–66%	$< 33\%$
maxIncrease/contractedPower	50kW	40kW	30kW
maxDurationMaxIncrease	240 min	120 min	60 min
maxDurationMaxDecrease	120 min	60 min	15 min
maxAdaptionTime	5 min	15 min	30 min
requestPeriod	block 3 hours	block 6 hours	block 9 hours

A. Monitoring Messages

In order to achieve the envisioned DR, it is essential for the EP to get certain information on the state of the participating DCs. Similarly, the DC might need DR-related data from the EP. In this section, messages are described that are used by the EP and DC to exchange state information.

1) GetPower – CurrentPower : For the case that the DC has no access to the smart meter that measures the power consumption of the DC, it has the possibility to request its current power consumption from the EP. The *GetPower* request needs to contain at least the `ID` of the DC and a `timestamp`. As a reply to the *GetPower* message, the EP creates a *CurrentPower* message which contains a `timestamp` and the `power` consumption (in Watt) of the DC for the time period specified by the timestamp. For the DC, it is essential to know its own power consumption so that during adaption it is able to verify whether the adapted power use complies with the request of the EP or not.

2) GetExpectedPower – ExpectedPower: The expected power consumption request message *GetExpectedPower* is used by the EP to ask for the expected power consumption of the DC for a certain future time period. Knowing the expected power consumption of the DC helps the EP to plan ahead and to be able to validate the power consumption adaption of the DC by comparing the actual power consumption of the DC to the predicted power consumption profile[2]. The *GetExpectedPower* message contains a `start_time` and an `end_time` indicating the beginning and the end of the time period for which the power consumption profile is requested. Furthermore, the request contains a field `resolution`[3] (in seconds) that defines the time between two power consumption values in the power profile. As a reply, the DC sends back to EP an *ExpectedPower* message which contains the `ID` of the DC, the `start_time`, the `end_time` and the `resolution`, similar to the requested message. Furthermore, it contains a vector of $\frac{\text{end\_time} - \text{start\_time}}{\text{resolution}}$ values indicating the power consumption (in Watt) profile of the DC for the requested time period.

3) GetExpectedGridLoad – ExpectedGridLoad: Similar to the *GetExpectedPower* message, the expected load curve request message *GetExpectedGridLoad* is used by the DC to ask from EP for its load curve of predicted surplus and shortage of energy in the grid for a given time period. Knowing the expected state of the grid helps the DC to adapt its long-term workload plan according to the state of the electricity grid. The *GetExpectedGridLoad* message contains a field for the requesting DC's `ID`. Furthermore, it contains fields for the `start_time` and `end_time` during which the expected load is requested. The field `resolution` (in seconds) indicates the desired resolution of the load profile (similar to the one of *GetExpectedPower*). The EP replies with an *ExpectedGridLoad* message. The fields in the *ExpectedGridLoad* message are the same as the ones for *ExpectedPower* message. However, positive values of the power vector indicate a surplus of power, whereas negative ones denote power shortage.

B. Adaption Messages

The described messages of this section are used when a power adaption is requested by the EP from the DCs.

1) AdaptReq: When the EP plans to compensate power surplus or shortage situations by using the DR approach, it broadcasts an *AdaptReq* message to all DCs participating in the DR program in order to obtain information on their adaption capabilities (see Algorithm 1). The *AdaptReq* message contains information about the `type` of the requested adaption, which is either "increase" or "decrease". If the type is "increase", the DC is requested to send information on the capabilities to increase its power consumption. If the type is "decrease", the DC is asked for information on its capabilities to reduce the power consumption. Furthermore, a field `duration` (in seconds), informs the DCs about the duration of the planned adaption phase. The field `start_time` gives the exact time when the adaption phase is planned to start. There are two possible responses that can be provided by the DC. Either they can acknowledge the request by sending at least one adaption profile or they can deny the request by sending a negative acknowledgement. The corresponding messages are described respectively in Sections V-B2 and V-B3.

2) AdaptACK: An *AdaptACK* message is sent to the EP by the DC if this latter is willing and able to adapt its power consumption for the requested time period. The *AdaptACK* message contains one or more adaption possibilities of the DC. The adaption possibilities are structured in the form of profiles with a unique `ID`. When the EP receives the *AdaptACK* message, it has the possibility to request a power consumption adaption of the DC by referring to one of the profiles in the *AdaptACK* message by using its ID (see Section V-B4). In more detail, an adaption profile has the following fields:

- `ID`: identifies the profile so that the EP can use it to refer to the corresponding profile.

---

[2] The DC uses prediction models to estimate the workload and the corresponding power consumption.

[3] For instance, every 5 minutes the power consumption is captured.

- The fields `type`, `duration` and `start_time` have the same meaning as the fields in the *AdaptReq* message.
- `power_to_adapt`: the amount of power adaption (in Watt) that corresponds to this profile.
- `rec_power`: the expected additional power consumption (in Watt) during the recovery phase.
- `rec_time`: the time when the recovery starts.
- `rec_duration`: the duration of the recovery phase in seconds.
- `validity`: the validity of this profile in seconds. If the EP requests the power adaption of a given profile, it has to request it before its validity expires.

3) AdaptNACK: An *AdaptNACK* message is sent to the EP if the DC is not willing or able to fulfil the adaption request. It contains only the `ID` of the DC and a field `type` with the value "NACK".

4) SelectProfile: The *SelectProfile* message is sent by the EP to the DC after the reception of an *AdaptACK* message. The *SelectProfile* message contains a field `profile_id`. Using this field, the EP refers to a profile in the *AdaptACK* message. When the DC receives the *SelectProfile* message, it adapts its power consumption according to the referred profile.

VI. Data Centre Scheduling Policy

As mentioned previously, whenever the EP faces a critical situation such as power surplus or shortage, it asks for the collaboration of DCs (that signed the GreenSDA) by sending them power adaption request (e.g. power to increase/decrease at a specific point and for a specified period of time). Once a DC receives such a request, it sends back all the possible adaption profiles (at least one) to EP.

In order to avoid all the DCs to react at the same time and hence create another critical situation in the near future, a limited and appropriate number of DCs needs to be contacted. In this section, we present a scheduling policy that tries to contact the DCs in a fair manner.

A. Problem Description

When the EP predicts a potential power shortage or surplus situations, it can estimate the amount of power to adapt (e.g. increase/decrease) as well as when and for how long the corresponding situation will occur.

Fig. 3(c) illustrates an example of a situation where an EP needs collaboration from DCs in order to avoid a shortage of 3 kW at the time instance $t_1 = 0$ for a duration of 60 minutes. This is depicted by the rectangular area in Fig. 3(c). Typically, the best solution is to fill in as much as possible this rectangular area with adaption capability blocks of DCs. Fig. 3(b) shows the adaption capabilities of three different DCs denoted respectively by green, orange and blue blocks. Note that each block corresponds to an adaption profile sent by the DC.

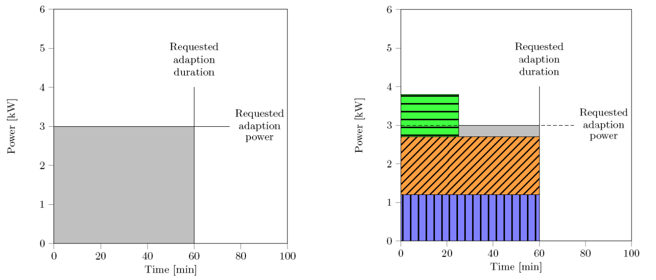

(a) Needed energy reduction (b) Possible energy adaption of DCs

Fig. 2. EP's required adaption power and duration

Thus, this requires a careful planning and selection of DCs from the EP's side in order not to create another critical situation once the current one is circumvented.

B. Algorithmic Overview

Algorithm 1 shows the necessary steps of a DC scheduling policy for power adaption purposes. We call such a policy "fair" since it tries to distribute the burden of collaboration among all DCs evenly.

Step 1 of Algorithm 1 guarantees that no adaption request is sent to a DC that has already reacted successfully to the previous requests and satisfied its promised collaboration as agreed in its GreenSDA contract. Step 2 ensures that the EP sends an adaption request (i.e. adaption start time, duration and power) to all (i.e. broadcast) the DCs to send back their possible adaption profiles with respect to the requested adaption start time, duration and power. Step 3 of Algorithm 1 is necessary to ensure not to select DCs whose recovery phase is running within the boundaries of the requested adaption phase. This is important to guarantee that no further shortages are created during the adaption. Step 4 ensures that the DCs are selected fairly enough based on the ratio of current to maximum requests achieved and the rectangular area is filled in appropriately. It is worthwhile to note that several iterations might be necessary within Step 4 of the algorithm until the most suitable scheduling plan can be found for the DCs. Step 5 guarantees that the DCs start their adaption based on the sent profile.

C. A Sample Scenario

In this section, we give a simple example just for clarification purposes of how Algorithm 1 selects DCs (detailed analysis will be provided in Section VII). In this example, we assume the followings:

- There is a shortage of 5 kW which is projected to occur at time instance $t_1 = 0$ for a duration of 60 minutes.
- Three DCs participate in the power adaption collaboration with the following ratio of *currentRequestsPerMonth* to *maxRequestsPerMonth*: 0.4 (DC_1), 0.8 (DC_2), and 0.4 (DC_3).

Algorithm 1 Data centre fair scheduling policy

```
REQUIRED:
$L$: List of all DCs that can collaborate with the EP
ENSURE:
DCs are selected in a fair manner in order to avoid shortage/surplus situations

BEGIN
1 - Add to the new list $L'$ all DCs found inside $L$ that haven't already achieved the maximum requests per month agreed in their GreenSDA (see Section IV-A) contract (e.g. currentRequestsPerMonth $\leq$ maxRequestsPerMonth).
2 - Send an adaption request (e.g. increase/decrease) to all DCs of $L'$ by specifying the adaption's start time (e.g. at 4:15 PM) and duration (e.g. 15 minutes)
    • Each DC upon receiving such a request sends back either the possible adaption profiles (e.g. the DC should send at least one adaption profile) or a negative acknowledgment
3 - Once all the replies of DCs are sent back:
    • For every negative acknowledgment or an empty adaption profile, increment by one the corresponding DC's currentRejectsPerMonth monitoring parameter
    • Eliminate those adaption profiles whose adaption duration is smaller than the requested adaption duration such that the corresponding profiles' recovery power is not zero
        - If the adaption profiles' power and duration fall outside the bounds of minimum and maximum adaption power and duration signed in the GreenSDA contract (Terms 1, 2, 3, and 4 of Section IV-A) or the adaption profiles' starting time is different then the requested starting time, then increment by one the corresponding DC's currentRejectsPerMonth monitoring parameter
4 - Based on the EP's requested adaption power and duration (e.g. rectangular area of Fig. 3(c)), choose the suitable DCs whose ratio of currentRequestsPerMonth to maxRequestsPerMonth is minimum
    • Add to the list $L''$ the corresponding DC
5 - Schedule accordingly based on the adaption profile blocks of DCs found inside the list $L''$
    • Send an acknowledgment for the adaption profile (identified by its unique number) back to DC
    • Increment by one the corresponding DC's currentRequestsPerMonth monitoring parameter
END
```

- DC_1 and DC_2 have the ability to reduce respectively 2500 W and 3000 W of their power consumption for a duration of 40 minutes without incurring any power for recovery. Note that such a reduction of power consumption in DCs is possible thanks to the workload consolidation (e.g. virtualisation) and turning off the unutilised servers.

Fig. 3 illustrates the different profiles that each DC sends once it has received an adaption request from EP.

We notice in Fig. 4 that the DCs 1 and 2 have both feasible adaption profiles after Step 3 of Algorithm 1 even though the profiles' adaption duration is smaller (both 40 minutes) than the requested adaption duration (60 minutes). The reason for this is that both profiles have a recovery power of zero and hence can be considered as valid profiles. Fig. 5 shows the result of applying Step 4 of the Algorithm 1. We see that in order to fairly distribute the adaption responsibilities among the DCs, the maximum adaption that can be achieved might not match to the required requested adaption.

VII. SIMULATION

To evaluate the potential to mitigate power shortages in the grid, a discrete-event simulation was implemented. It is used to simulate the power grid of E.ON, a major German EP.[4] The power grid of E.ON covers about 39% of Germany.

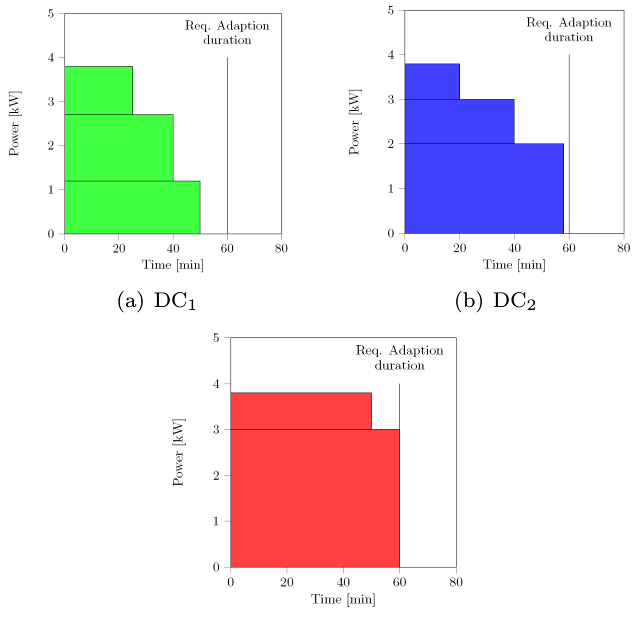

(a) DC_1 (b) DC_2 (c) DC_3

Fig. 3. The different adaption profiles of DCs

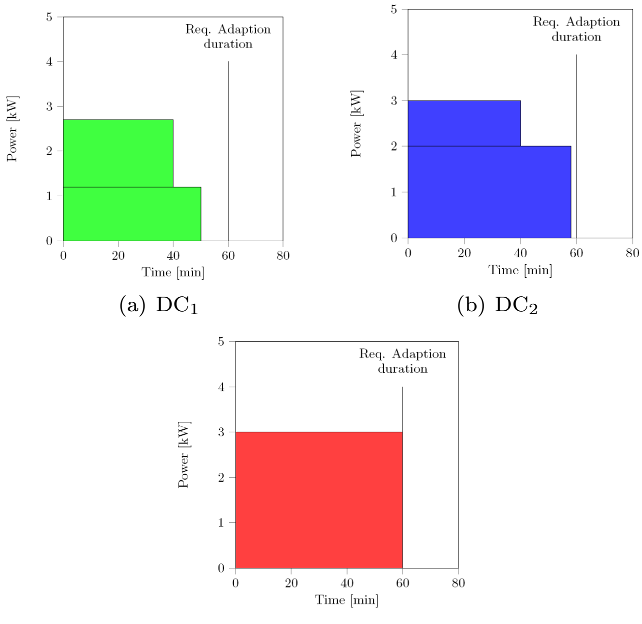

(a) DC_1 (b) DC_2 (c) DC_3

Fig. 4. Remaining profiles of DCs after Step 3 of the Algorithm 1

A. Simulation Setup

The power grid state was evaluated for a time period of one year, in steps of one minute. As the temporal resolution of the available grid load data is 15 minutes,

---

[4] The power grid data is available online http://www.eon-netz.com/pages/ehn_de/Veroeffentlichungen/Netzkennzahlen/Lastverlauf/index.htm

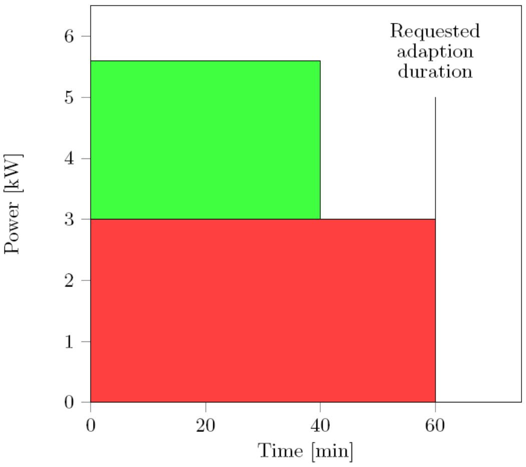

Fig. 5. The selected DC's for adaption

the missing data was interpolated linearly. The statistics regarding the distribution of DC sizes and corresponding power consumption in Germany were taken from [6], [7]. These reports cover statistics on the number and power demand of German DCs, and categorise them into five sizes:

- Rack
- Server room
- Small DC
- Medium DC
- Large DC

For the simulation, the sizes 'Rack' and 'Server room' were ignored. Adjusting the numbers obtained from [6], [7] this results in the DC size distribution listed in Table II.

DC size	Fraction of total DC number
Small	76%
Medium	22%
Large	2%

TABLE II. DC SIZE DISTRIBUTION USED IN THE SIMULATION

The respective power consumptions for large, medium and small DCs are 5700 kW, 550 kW and 105 kW. Both statistics for the power grid and DCs cover the year 2011. The following assumptions were made for the simulation:

1) The low adaption profile of the DCs does not need recovery power. Medium and large adaption profiles do.
2) The medium adaption profile offered by DCs is derived by calculating the mean of low and high adaption profiles (both in amount and duration).
3) The grid will be regarded as in a state of shortage, whenever the total load of the grid is within 2% of the highest value of the entire year.
4) Only the following GreenSDA parameters were considered in the simulation:
    - *maxRejectsInSuccession*
    - *maxRejectsPerMonth*
    - *maxRequestsPerMonth*
    - *maxDecrease*
    - *minDecrease*
    - *maxDurationMinDecrease*
    - *minDurationMinDecrease*
    - *maxDurationMaxDecrease*
    - *minDurationMaxDecrease*
5) The simulation will only deal with power shortages, surpluses will not be considered (as the identification of surpluses requires additional data).
6) The EP is assumed to be able to accurately predict its grid state 10 hours into the future.

The simulated DCs were equipped with randomly generated GreenSDAs, such that the individual values were taken from discrete uniform distributions with boundaries listed in Table III.

Parameter	Range
minDecrease	[5, 15]
minDurationMinDecrease	[10, 120]
maxDurationMinDecrease	[120, 600]
maxDecrease	[15, 66]
minDurationMaxDecrease	[5, 15]
maxDurationMaxDecrease	[15, 120]
maxRejectsPerMonth	[1, 2]
maxRejectsInSuccession	[1, 2]
maxRequestsPerMonth	[1, 4]

TABLE III. PARAMETERS USED IN THE SIMULATION.

Three different scenarios were simulated: *full*, *medium* and *low* cooperation of the DCs. In the full cooperation case, all DC always offer their maximum adaption capabilities - both in amount and duration. No DC rejects requests even if the maxRequestsPerMonth parameter of its GreenSDA has already been reached. In the medium cooperation case, DC still try to serve all requests, however they do not necessarily offer all possible adaption profiles. Also, they only accept requests until their maxRequestsPerMonth have been used up. In the low cooperation case, it is more unlikely that DCs offer all profiles, and additionally they may also reject requests in case the maxRejectsPerMonth and maxRejectsInSuccession parameters are not exceeded.

For all three scenarios, different numbers of DCs were tested regarding their ability to collectively mitigate the occurance of power shortages in an electricity grid.

Additionally, the DC scheduling algorithm (see Algorithm 1) was tested for the load balancing quality exhibited. The average GreenSDA load (i.e. the mean of the ratio of actually served requests to the maximum possible requests per year among all DCs) and the standard deviation of this value regarding all DCs was calculated for the different scenarios.

B. Simulation Results

First, the results regarding different cooperation levels of DCs will be discussed. Fig. 6 shows that the number of DCs required to cope with large power shortages is largely dependent on their degree of cooperation. With full cooperation, a rising number of participating DCs quickly leads to a converging result of complete recovery from the shortage. Due to the large number of participants, the random creation of GreenSDAs does not significantly reduce the results confidence intervals. For the medium cooperation case, a higher number of participating DCs is needed, mainly caused by the limit of monthly requests a DC will serve. However, this level of cooperation is still sufficient to deal with peaks to a large extend. In contrast, when looking at the low cooperation scenario, it becomes apparent that DCs just offering the bare minimum to fulfill their GreenSDA severely diminish the effectiveness of our approach. In summary, the simulation supports the expected results: higher cooperation of the DCs reduces the required number of participants to cope with power shortages. Even a medium cooperation behaviour of DCs

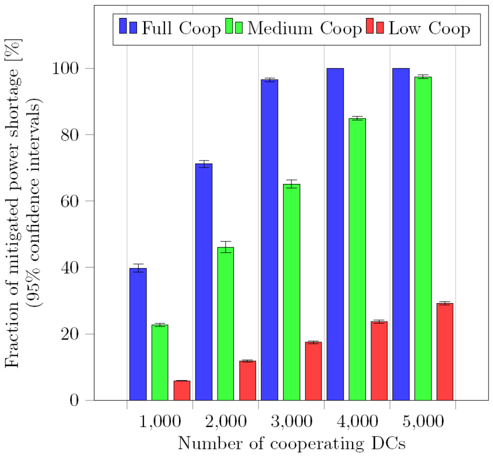

Fig. 6. Capability of different numbers of DCs to mitigate power shortages

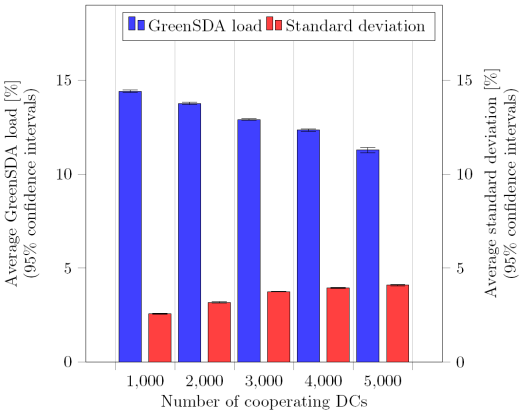

Fig. 7. Average DC GreenSDA load and standard deviation

is sufficient to cope with large amounts of power shortages. Regarding the quality of load balancing among the DCs, the medium cooperation scenario was evaluated in greater detail. Fig. 7 shows the results of the simulation runs.

The figure shows two criteria regarding the load balancing: The average GreenSDA load and its average standard deviation. As expected, the GreenSDA load drops with an increasing number of participating DCs, as the power adaption requests can be fulfilled by fewer DCs. However, as the power shortages are concentrated during winter months, the observed decrease may appear less strong than expected.

The standard deviation is fairly stable and increases only slightly with rising number of participating DCs. This increase is caused by not all DCs being fully loaded anymore, therefore slight imbalances in GreenSDA load occur. However, overall a high quality load balancing is achieved with high confidence.

VIII. CONCLUSION

For the EP, the main challenge is to maintain the stability of the power grid. A main requirement for stability is to match power demand and supply. However, there are certain situations where the demand and supply do not necessarily match. To circumvent such critical situations, the EP may rely on the collaboration and cooperation of its customers. To this regard, DCs play an important role due to their high energy use as well as the automated IT infrastructure. However, the major drawback of today's electricity tariffs between EP and its customers is the lack of flexibility. Consequently, incorporating power adaption collaboration between EP and DCs becomes challenging.

In this paper[5], we introduced new Green Supply Demand Agreements in order to leverage power adaption collaboration between the EP and DCs. Based on the proposed agreement terms, we defined reward and penalty schemes as well as communication requirements for signalling. Moreover, we provided a scheduling policy that guides the EP in selecting DCs to ask for power adaption. Finally, we set up a simulation environment by analysing the grid load of a major German EP, E.ON. The simulation confirmed the expected results: higher cooperation of the DCs reduces the required number of participants to cope with power shortages. Even a medium cooperation behaviour of DCs is sufficient to cope with large amounts of power shortages.

IX. ACKNOWLEDGEMENT

The research leading to these results was supported by the European Community's 7^{th} Framework Programme in the context of the All4Green project.

REFERENCES

[1] *Demand Response in Electricity Markets: An Overview*, June 2007.

[2] H. Aalami, M. P. Moghaddam, and G. Yousefi. Demand response modeling considering interruptible/curtailable loads and capacity market programs. *Applied Energy*, 87(1):243 – 250, 2010.

[3] R. Basmadjian, H. D. Meer, R. Lent, and G. Giuliani. Cloud computing and its interest in saving energy: the use case of a private cloud. *Journal of Cloud Computing: Advances, Systems and Applications*, 1(5), 2012.

[4] C. Dupont, G. Giuliani, F. Hermenier, T. Schulze, and A. Somov. An energy aware framework for virtual machine placement in cloud federated data centres. In *Proceedings of the 3rd International Conference on Future Energy Systems (e-Energy 2012)*. ACM, 2012.

[5] G. Ghatikar, M. A. Piette, S. Fujita, A. McKane, J. H. Dudley, and A. Radspieler. Demand response and open automated demand response opportunities for data centers. Technical report, LBNL, 2010.

[6] R. Hintemann and K. Fichter. Materialbestand der Rechenzentren in Deutschland - Eine Bestandsaufnahme zur Ermittlung von Ressourcen- und Energieeinsatz. *Texte*, 55/2010, 2010.

[7] R. Hintemann and K. Fichter. Energieverbrauch und Energiekosten von Servern und Rechenzentren in Deutschland - Aktuelle Trends und Einsparpotenziale bis 2015. 2012.

[8] OASIS Committee Specification 02. Energy interoperation version 1.0, February 2012.

[9] M. A. Piette, G. Ghatikar, S. Kiliccote, E. Koch, D. Hennage, P. Palensky, and C. McParland. Open automated demand response communications specification (version 1.0). California Energy Commission, 2009.

---

[5] The proposed approach has been implemented within the context of All4Green project, and the trial results are expected to be obtained by end of June 2013