Energy Supply Aware Power Planning for Flexible Loads

Florian Niedermeier
University of Passau
Innstraße 43
Passau, Germany
florian.niedermeier@uni-passau.de

Fiodar Kazhamiaka
University of Waterloo
200 University Avenue West
Waterloo, Ontario, Canada
N2L 3G1
fkazhami@uwaterloo.ca

Hermann de Meer
University of Passau
Innstraße 43
Passau, Germany
hermann.demeer@uni-passau.de

ABSTRACT

Increasing the use of renewable energy is considered a viable way of reducing carbon intensive power generation. However, a power grid running on high amounts of renewable energy has to deal with the limited controllability and higher volatility of power sources like wind or solar. In this work, we propose to use demand side management to deal with varying amounts of renewable power feed-in via the use of power plans, i.e. instructions passed to large energy consumers that specify how they should try to spread out their energy use over a day. We argue that a separation of power planning and implementation of technical measures to schedule loads to follow the plan would alleviate some of the problems faced by an integrated planning-scheduling approach, as these processes are governed by different entities who may be unwilling to disclose all required information to each other. As a proof-of-concept, we propose and analyze a quadratic programming approach to maximizing the fraction of renewable energy being used while not overburdening the consumer with a power plan that is difficult to follow.

CCS Concepts

•Mathematics of computing → Quadratic programming; •Software and its engineering → Power management;

Keywords

Power Planning, Quadratic Programming, Flexibility

1. INTRODUCTION

With urbanization progressing, there is a legitimate demand for large energy consumers in cities to be powered from renewable sources. However, as two major sources of renewable energy, solar and wind, are exhibiting volatile availability patterns in both long and short terms, in order to maintain power grid generation/demand equilibrium a mechanism is required to compensate for these fluctuations. There are two main ways of coping with fluctuations from the production side: curtail renewable energy generation when it is too high, or lower the utilization of fuel-based power plants to allow for high penetration of renewables. Both approaches have significant down-sides: curtailment of renewable sources is effectively throwing away 'free' power since the cost for renewable generation dominated by equipment cost which is paid up-front, while low utilization of fuel-based power plants greatly lowers their efficiency [12]. An alternative that avoids the aforementioned downsides is to take advantage of flexible loads via demand side management [7, 8]. One effect of the volatile feed-in of renewable energy to the grid is a varying percentage of green energy inside the energy mix delivered to consumers. Using power at times with high renewable concentration is an ecologically desirable goal, and the availability of predictions for the amount of renewable energy in the grid [11] suggests the opportunity for the consumer to plan their energy consumption accordingly. If a consumer were to shape their load to correlate with the grid penetration of renewables, this would aid in reducing the amount of reserve power generators that have to run at low utilization and reduce the amount of curtailment to renewable sources in order to achieve supply-demand balancing of power. At the same time, the flexibility constraints of the consumer must be taken into account.

To increase the share of renewable energy while at the same time retaining the ability to deliver uninterrupted service, the consumer can be provided with instructions on when to use which amount of power. We call these instructions a *power plan*.

The contributions of this work are the following:

1. We introduce the concept of power planning and argue why it is reasonable to strictly separate it from measures that implement a power plan.
2. We derive a widely applicable cost function which is capable of model energy consumer constraints regarding power flexibility. The cost function may be tailored to reflect the peculiarities of different consumer constraints by adapting its three parameters.
3. We present a quadratic programming based approach to optimal power planning for renewable energy intake that balances the objectives of the producer and consumer of energy. We test this approach using real data, and conduct a sensitivity analysis to model different consumer flexibility across multiple parameters.

The remainder of this paper is structured as follows: Section 2 summarizes existing work related to the topic. In Section 3, we discuss the motivation of creating a separate power planner, as well as consumer concerns and assump-

E2DC, June 21-24 2016, Waterloo, ON, Canada

ISBN 978-1-4503-4421-0/16/06...$15.00

DOI: http://dx.doi.org/10.1145/2940679.2940681

tions. Section 4 presents the formulation of our optimization problem and modelling of consumer flexibility. Section 5 presents the our analysis, and possible applications are discussed in Section 6. The paper is concluded in Section 7.

2. RELATED WORK

There has been a lot of recent work focused on scheduling flexible loads so that high energy usage corresponds to periods of high renewable energy generation. Data centers, with their high electricity demand and relatively flexible workloads, have been used as a model application for this research. In [6], Liu et al. present a linear programming based optimization for virtualized data centers (DCs). Their solution is specifically tailored towards DCs in a combined power planning and scheduling approach, and they propose a program to generate server workload schedules that maximizes the intake of renewables by a DC. Similarly, Rao et al. [13] solve a mixed-integer programming problem to minimize energy cost for DC. Zhu et al. [14] propose a rolling-horizon scheduling architecture for virtualized cloud DCs. Again, their approach integrates scheduling and power planning and is performed in an online manner, recalculating each time a task arrives into the system. Aksanli et al. [1] propose a job scheduler that takes into account green energy predictions while reducing the number of job cancellations due to fluctuations in green energy.

All the approaches mentioned differ from the one proposed in this work as they use an integrated system tailored to a specific use case, while in this work power planning is decoupled from the actual technical implementation (schedule) on the consumer side, allowing for lower computational complexity and wider applicability. Energy planning, often using multicriteria decision-making [9, 10], differs from our approach in two ways: first, it is traditionally oriented towards the power generation side. Second, its time frame is much longer, guiding a more strategic development rather than intra-day optimization.

3. SEPARATION OF CONCERNS: A DEDICATED POWER PLANNER

The power planner is a component in a larger integrated system, developed in the course of the EU FP7 project DC4Cities[1], which is briefly explained in the following. The DC4Cities system consists of three major components: data gathering, processing, and actuation. Without going into too much detail, data gathering is concerned with obtaining information on renewable energy availability and power grid energy mix details. Using this data, the processing stage calculates how much power should be assigned to each service present in the system. Finally, actuation is responsible for enacting the changes decided in the processing stage.

The power planner focused in this work is a subsystem at the very beginning of the processing stage, calculating when to use which amount of power. To this end, it utilizes the information on renewable power availability and grid energy mix. Upon completion, the power plan (i.e., power to use in future) is handed to the load to implement the according measures to adapt its power demand. Figure 1 shows the overall high level DC4Cities architecture.

While there are already several approaches to combined power planning and its implementation (cf. Section 2), we

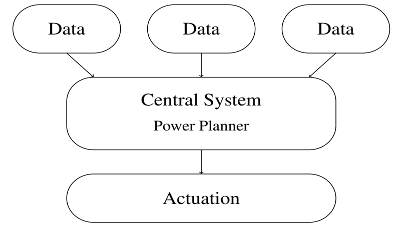

Figure 1: The high level DC4Cities architecture.

believe there are strong arguments for why it is useful to decouple power planning from the actual implementation of power adaptation mechanisms on consumer side:

1. Separate entities may be responsible for power planning and for the technical implementation of measures to reach the desired power demand. Consider a "Smart Grid" scenario where the goal is to increase the intake of renewable energy by an electrical load. The grid operator has a clear view on the mix of renewable energy in the grid, as well as other constraints and objectives that should be considered to ensure the stability of the grid (some of which could be restricted/sensitive information) when designing a power plan. The load operator has a clear view of the flexibility they have in scheduling work, but this information may be confidential. These two entities would have to either exchange sensitive information or use a trusted third party in order to use an integrated power planning tool, neither of which are attractive options.
2. Not all scheduling constraints may be expressed in a way that is suitable to be included in a mathematical program, which limits the applicability of an integrated planning and scheduling approach that uses this tool. Separating planning from scheduling allows for a dedicated first computation of an appropriate power plan using mathematical optimization, and a second stage of implementation with a different set scheduling tools that can deal with complex requirements.

The goal of a power plan from the grid operator's point of view is to suggest a power profile to a flexible load to increase its uptake of renewable energy and thereby reduce the magnitude of renewable energy fluctuations that have to be met with alternative sources. At the same time, the power plan is only useful if it is followed by the consumer, and therefore it must take into account constraints on consumer load and flexibility. The consumer may have different levels of flexibility with respect to changes in the power plan over time, oscillations in the power plan, and power ramps, to name a few. Therefore, we define the *power planning problem* as follows: Given information on renewable energy availability and the cost function of the consumer, derive a power plan that takes into account both maximizing the uptake of renewable energy and the costs associated with deviating from the consumers baseline power demand.

We cannot divide any integrated approach into two parts (planning and scheduling) without first considering that the power plan needs to have certain qualities to help ensure that there is a high chance the scheduler can implement it without large deviations. With an integrated power planner and scheduler, this is avoided because it cannot propose a plan without also proposing the schedule. We can imagine a scenario where a power plan maximizes the objective of increased renewable penetration, but oscillates due to the fluctuations in renewable generation, which could then make it

[1]http://www.dc4cities.eu/

difficult to derive a feasible schedule that matches the plan. In a data center setting, for example, sharp fluctuations in a power plan would mean the frequent shutting down and powering up of servers in order to achieve a decrease in power consumption, which is harmful to equipment lifetime [5]. The power planner should to some extent be aware of what a reasonable power plan looks like in order to avoid making it difficult for the scheduler to do its job. Our solution to this issue is to incorporate common consumer demands into a general objective function, with weights that can be tuned to match the application's requirements.

3.1 Consumer Concerns

To be applicable to a wide variety of loads, we model the cost of a power plan as penalties on three possibly challenging characteristics of the power plan:

1. Deviation from unaltered baseline power
2. High frequency power changes
3. Low lead time power changes

Each of these cost terms may be weighted to reflect the particularities of the load. Details on the choice of weights are discussed in Section 4.3. An additional concern is that the load has to remain capable of providing the same service as in the baseline scenario, i.e. the overall energy suggested by a new power plan should remain unchanged from the original power plan for the same time horizon. The authors are aware that overall energy demand may not always be proportional to service quality, however for many types of loads (e.g., EV battery charging) it is well suited as a measure of service delivered.

3.2 Assumptions

In our work, we assume a power planner is used as a means to increase the uptake of renewable energy while allowing the consumer to express their flexibility to change their planned consumption. It is assumed that the load has the capability to shift some or all of its electrical power demand in time to a certain extent. The load is connected to at least one power grid with an energy mix containing a time-varying amount of renewable energy in the mix. The loads could also be partially met with locally generated renewable sources.

4. METHODOLOGY AND FORMULATION

To provide a constant guide on which amount of power to consume, the power planning must be provided in a continuous manner. To this end, power planning may be performed in a rolling horizon approach, similar to the work of Li et al. [4], who use rolling horizon to manage production capacity. In contrast to our approach however, they perform an integrated planning-scheduling approach as opposed to the separated approach more commonly seen in industry [3]. A rolling horizon approach requires the power planner to generate a new power plan at each time slot over a sliding window. The final power plan consists of each first time slot of the calculated power plans as the window moves through time. In our examples, we consider a window length of 24 hours, where we could expect accurate forecasts and take advantage of the 24 hour period typically exhibited by renewable energy sources. A rolling horizon approach has several advantages compared to following a power plan calculated once every 24 hours:

1. Recalculating the power plan will exploit the higher accuracy of forecasts over a shorter time horizon.

2. Errors in the forecast of required power at consumer side can be dealt with by assigning additional power.

In Section 5.1, power planning in a rolling horizon configuration is analyzed.

4.1 Formulation

We formulate the problem of creating a power plan as a multi-objective optimization problem where we optimize the power planned in each time slot i (of equal length) in $[1, T]$ with T being the planning horizon. Descriptions of the parameters and variables can be found in Table 1.

Table 1: Notation

Name	*Description (units)*
Parameters	
$Pold_i$	Power allocated to time slot i in previous power plan (Watts)
$RenPercent_i$	Forecast of renewable percent in the grid in time slot i (%)
$LocalRen_i$	Forecast of locally generated renewable power in time slot i (Watts)
$Pmin_i$	Minimum power required by the consumer in time slot i (Watts)
$Pmax_i$	Maximum power the consumer is able to consume in time slot i (Watts)
α, β, γ	Weight parameters for cost function
T	Number of continuous time slots considered
Variables	
$Pnew_i$	Power allocated to time slot i (Watts)
$Cost_i$	The cost of implementing a new power plan; an offset to the renewable grid energy being used in time slot i

Given $(Pold_i)$, $(RenPercent_i)$, $(LocalRen_i)$, T, $Pmin_i$, and $Pmax_i$:

$$\max_{Pnew_i} \sum_{i=1}^{T} \Big[(Pnew_i - LocalRen_i)RenPercent_i - Cost_i\Big] \quad (1)$$

subject to

$$Pmin_i \leq Pnew_i \leq Pmax_i \qquad \forall i \in [1, T], \text{ and} \quad (2)$$

$$LocalRen_i \leq Pnew_i \qquad \forall i \in [1, T], \text{ and} \quad (3)$$

$$\sum_{i=1}^{T} Pold_i = \sum_{i=1}^{T} Pnew_i. \quad (4)$$

The objective function maximizes the amount of energy that the load uses from renewable sources that are coming from the grid, i.e., the difference between the allocated power and forecasted local renewable generation. It is also affected by a cost function, which is discussed in Section 4.2. Constraint (2) specifies that the power plan cannot give values outside the feasible range that the load is capable of using. Constraint (3) specifies that all of the local renewable generation must be used. This constraint is reasonable when we consider typical large energy consumers, who would rarely have a local renewable source that would generate more than the load during regular operation, and greatly simplifies the formulation. Constraint (4) specifies that the total amount of energy allocated in the new power plan is the same as in the old power plan. Additional application-dependent constraints could be added to the formulation as necessary.

This formulation can be used to implement a rolling horizon power planner, making use of updated forecasts together with the previously computed power plan to compute a new plan.

4.2 Cost Function

In addition to the deviation from the baseline power demand, the cost function $Cost_i$ should capture the consumer concerns outlined in Section 3.1: 1) It is preferable to avoid spikes in the power plan, 2) The power plan for upcoming time slots is harder to modify than the power plan for temporally distant time slots, 3) Deviations and ramps often have a cost that is non-linear with respect to their magnitude. The following cost function reflects these measures of flexibility:

$$Cost_i = \gamma \frac{(|Pold_i - Pnew_i|)^2}{1+\alpha i} + \beta(|Pnew_i - Pnew_{i+1}|)^2 \qquad (5)$$

Here, α is a linear scaling constant that expresses at what rate we relax the cost of deviations from the old power plan as we look further into the future, and β is the cost attributed to having sharp increases in the power plan. The second term is set to 0 when $i = T$, since there is likely no existing plan for the last time step in the planning window. Through trial and error, we discovered that squaring the cost of deviations and ramps leads to a power plan with desirable characteristics, such as smoothness. The absolute value term can be linearized, but the cost function has two quadratic terms and is not necessarily convex. To simplify the function, we combine all parts of the cost function into one quadratic term, and after linearizing the absolute values we get the following form:

$$Cost_i = \left(\gamma \frac{X_{1i} + X_{2i}}{1+\alpha i} + \beta(X_{3i} + X_{4i})\right)^2, \forall i \in [1, T] \qquad (6)$$

$$X_{1i} - X_{2i} = Pold_i - Pnew_i, \qquad \forall i \in [1, T] \qquad (7)$$

$$X_{3i} - X_{4i} = Pnew_i - Pnew_{i+1}, \qquad \forall i \in [1, T-1] \qquad (8)$$

$$X_{3T} + X_{4T} = 0 \qquad (9)$$

$$0 \leq X_{1i}, X_{2i}, X_{3i}, X_{4i}, \qquad \forall i \in [1, T] \qquad (10)$$

where $X_{1i}, \ldots, X_{4i}$ are variables introduced to linearize the absolute values. Note that incorporating this cost function into our formulation results in a convex quadratic problem.

4.3 Interpretation of γ, α, and β

The objective function tries to model the desirable properties of a power plan. First and foremost, the goal is to try and plan for increased usage of renewable generation. However, the plan cannot be made from scratch, and large changes to the power plan should be avoided. The parameter γ is the cost weight that we assign to the difference (in kW) between the old and new power plans.

The flexibility of the power plan is not constant over time; it is much easier to plan and prepare changes for 20 hours in the future, rather than 1 hour in the future. The parameter α can be interpreted as a decay rate of the cost assigned to changes in the power plan. A high value of α means that the the cost of changes for the last time slots is much lower than the cost of changes for the first time slots.

The power plan is much easier to follow if it does not have steep ramps. Oscillations in the renewable energy generation predictions (for example, predicted intermittent cloud cover) could lead to a power plan that maximizes usage of renewable generation by following these oscillations. The parameter β is the cost weight we assign to the difference in power between adjacent time slots in the new power plan.

5. ANALYSIS

In the following, we analyze the impact of changing parameters α, β and γ on the calculation of a single power plan. Additionally, we provide an example of deriving a power plan in a rolling horizon configuration. We assume a load with a minimum power demand of 50 kW, a maximum power demand of 500 kW, and a base power plan of 225 kW of consumption for the entire planning window. We fix a planning horizon of 24 hours. Service scheduling or deadlines are not considered yet, i.e. only the amount of power to be consumed in different time slots is calculated. The authors are aware that in a later stage, the power values calculated have to be reached by adapting the load accordingly. Figures 2, 3 and 4 show examples of output power plans obtained by varying one parameter of Eq. (6). The fraction of renewable energy in the energy mix (REN%) is depicted as green area. Data on REN% is taken from both German (2015) and Spanish (2013) power grid data. For each figure, quantitative results are provided in a table, showing the fraction of renewable energy in the calculated power plan. Additionally, the maximum required flexibility is given, i.e. the maximum amount of power the load has to deviate from a constant power demand of 225 kW.

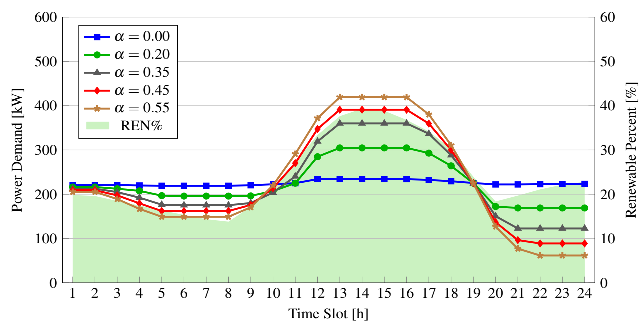

Figure 2: Results for $\beta = 0.03$, $\gamma = 0.08$, varying α

Parameter α gives a discount on changes that occur in time slots further in the future. Figure 2 shows a low deviation from the baseline in early time slots, while there is increasingly strong adaptation in later time slots. It is noteworthy that even though REN% is lower in early time slots, stronger deviations from the baseline occur in later time slots (discount from α outweighs high adaptation amount costs in late time slots).

Table 2: Results for $\beta = 0.03$, $\gamma = 0.08$, varying α

α	Renewable energy from grid in power plan [%]	Maximum flexibility required [kW]
0.00	24.16	9.25
0.20	25.60	79.68
0.35	26.72	134.99
0.45	27.40	165.93
0.55	28.00	194.34

Figure 3 shows the effects of different costs for high frequency adaptations ($\beta(Pnew_i - Pnew_{i+1})$). A higher value of β has two major effects: first, the ramping times of the calculated power values increase, leaving the consumer with more time to gradually adapt its power demand. As a result, adaptation to changes in REN% is slower than the changes in REN% and, hence, the overall use of renewables is lim-

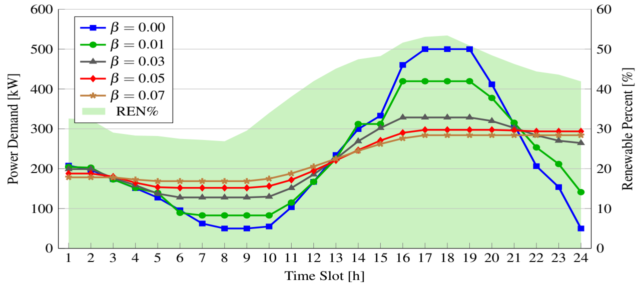

Figure 3: Results for $\alpha = 0.35$, $\gamma = 0.08$, varying β

ited as the power plan does not correlate as strongly to the valleys and peaks in REN%.

Table 3: Results for $\alpha = 0.35$, $\gamma = 0.08$, varying β

β	Renewable energy from grid in power plan [%]	Maximum flexibility required [kW]
0.00	44.97	275.00
0.01	44.07	194.48
0.03	42.64	103.55
0.05	41.93	73.09
0.07	41.53	59.01

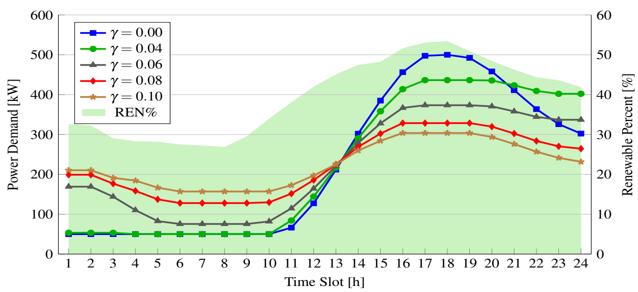

Figure 4: Results for $\alpha = 0.35$, $\beta = 0.03$, varying γ

Finally, the result of changes to γ is depicted in Figure 4. Increasing γ will result in rising costs to deviations from the baseline power demand. As expected, higher values will result in a smoother power plan, however limiting the increase of REN% in total energy.

Table 4: Results for $\alpha = 0.35$, $\beta = 0.03$, varying γ

γ	Renewable energy from grid in power plan [%]	Maximum flexibility required [kW]
0.00	46.82	275.00
0.04	46.32	211.40
0.06	44.34	149.50
0.08	42.64	103.55
0.10	41.72	78.66

Figure 5 shows a series of pareto curve, which was calculated by solving a modified optimization problem with different fixed values of renewable energy usage and getting the lowest cost of the power plan that can results in each fixed renewable energy usage. Each curve corresponds to a different set of cost parameters. A pareto curve can be used to guide the producers and consumers in choosing a power plan with a trade-off that is acceptable to both parties.

5.1 Rolling Horizon

The result of the rolling horizon test case depicted in Figure 6. This example initiates the power plan to be a flat line at 225 kW, and recalculates the power plan every hour. Intermediate results for several individual runs of the power planner (thin lines) and the resulting final power plan (marked with red circles) are displayed. In this example, the parameter α in the cost function is set high enough to prevent large changes to time slots in the near future. In our example, we did not change the forecasted REN% over time, though in practice the forecasts would improve as more data is gathered about upcoming time slots. The power planner would be able to adjust to match new forecasts.

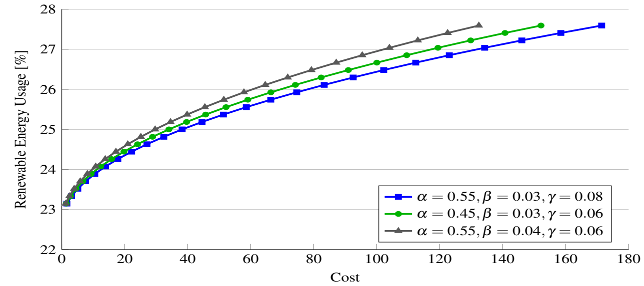

Figure 5: Pareto curve for cost and renewable energy consumption.

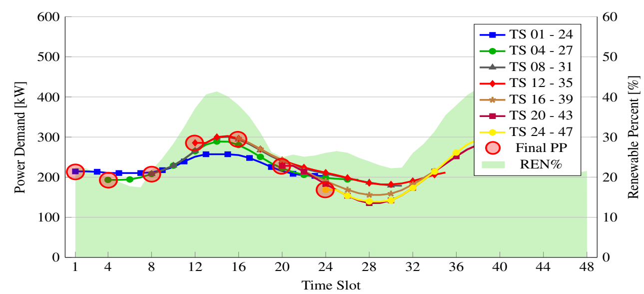

Figure 6: Result of successive power plan calculations to form a rolling horizon power plan.

6. USE CASES

In the following, two possible use cases for power planning are discussed. The use case specific influence on cost function parameters and constraints are discussed.

6.1 Data centers

Data centers (DC) are major consumers of electrical power. Depending on the type of services offered by a DC, a certain amount of work may be shifted in time. In comparison to other flexible loads, DCs are already well suited to implement power plans as there are automation frameworks usually already in place. These may be used to automate the interaction with power planning, e.g. by using the OpenADR protocol[2]. To be well suited to adapt to a power plan, a DC needs to be capable of providing flexibility in power demand. This can be achieved by rescheduling jobs in a way that overall IT utilization (and therefore power demand) adapts to the demands of the power plan. The amount of workload which may be rescheduled is limited by the ratio of interactive (not delay tolerant) to non-interactive (delay tolerant) workload and may vary strongly between different DCs. However, if details are unknown, a ratio of 1:1 seems a

[2]http://www.openadr.org/

valid assumption following a discussion in, e.g., [6]. It should be noted that for DCs, not only the IT power demand but also related infrastructure (especially cooling) contribute a significant amount to the overall power demand. This additional power demand will also be decreased by a reduced IT load (for details, see e.g. [2]). The cost function parameter α is increased in case the DC requires planning certainty for the time slots in the near future. β would be set high enough to limit high frequency oscillations between adjacent time slots, which could lead to repeated shutdown and start of servers. Finally, the value of γ determines how strong the overall deviation from an unaltered power demand is penalized. γ is set according to how many shiftable jobs are available.

6.2 Electric vehicles

Electric vehicles are assumed to be of increasing importance in the future. With a rising market penetration, the cumulative power draw of large numbers of electric vehicles (EV) will be significant. A power planning process is essential to fulfill EV users demands regarding their battery state of charge (SoC) requirements while using as much renewable energy as possible. The proposed power planning is suitable for use at an EV charging station to plan charging with minor adaptations:

1. Parameters α, β, γ will mainly reflect constraints regarding the state of health of the battery, where a prolonged high charging current has the potential to damage the EV battery.
2. Constraints are needed to set the maximum charging current and minimum SoC at certain times.
3. In certain cases such as vehicle-to-grid applications, it may also be feasible to discharge the battery to support the power grid. Therefore, negative values could be allowed in the power plan.

It is noteworthy that in order to reliably provide an EV user with the desired SoC at the end of the charging interval, the power plan duration should match the estimated time the EV is connected to the grid. As previously discussed, the power plan will make sure to deliver the same total amount of energy as in the baseline scenario, so the SoC will not be influenced. In case the EV user wants to charge quickly, a baseline with high power demand in the first time slots may be used in conjunction with a high α value to promote quick charging in the first time slots.

7. CONCLUSIONS AND FUTURE WORK

In this paper we presented a general quadratic programming based power planner that can be tailored to many applications, for the purpose of maximizing the fraction of renewable energy in a load while considering the challenges on consumer side to adapt its power demand. We argue that it is reasonable to split power planning from concrete power adaptation implementation on consumer side, and conducted a sensitivity analysis to show how the power plan can be tailored to cover three major facets of consumer load flexibility.

An extension of this work could consider additional constraints in power plan flexibility in order to capture an even wider set of applications. It would also be worthwhile to explore the mechanisms for incentives that motivate the consumer to implement a power plan.

8. ACKNOWLEDGEMENTS

This work was partially carried out within the European Project DC4Cities (FP7-ICT-2013.6.2).

9. REFERENCES

[1] B. Aksanli, J. Venkatesh, L. Zhang, and T. Rosing. Utilizing green energy prediction to schedule mixed batch and service jobs in data centers. *ACM SIGOPS Operating Systems Review*, 45(3):53–57, 2012.

[2] C. Belady, A. Rawson, J. Pfleuger, and T. Cader. Green grid data center power efficiency metrics: PUE and DCiE. Technical report, Green Grid, 2008.

[3] J. Kallrath. Planning and scheduling in the process industry. *OR spectrum*, 24(3):219–250, 2002.

[4] Z. Li and M. G. Ierapetritou. Rolling horizon based planning and scheduling integration with production capacity consideration. *Chemical Engineering Science*, 65(22):5887 – 5900, 2010.

[5] M. Lin, A. Wierman, L. L. Andrew, and E. Thereska. Dynamic right-sizing for power-proportional data centers. *IEEE/ACM Transactions on Networking (TON)*, 21(5):1378–1391, 2013.

[6] Z. Liu, Y. Chen, C. Bash, A. Wierman, D. Gmach, Z. Wang, M. Marwah, and C. Hyser. Renewable and cooling aware workload management for sustainable data centers. In *ACM SIGMETRICS Performance Evaluation Review*, volume 40, pages 175–186, 2012.

[7] P. Palensky and D. Dietrich. Demand side management: Demand response, intelligent energy systems, and smart loads. *Industrial Informatics, IEEE Transactions on*, 7(3):381–388, 2011.

[8] M. Paulus and F. Borggrefe. The potential of demand-side management in energy-intensive industries for electricity markets in germany. *Applied Energy*, 88(2):432–441, 2011.

[9] S. Pohekar and M. Ramachandran. Application of multi-criteria decision making to sustainable energy planning-a review. *Renewable and Sustainable Energy Reviews*, 8(4):365 – 381, 2004.

[10] H. Polatidis, D. A. Haralambopoulos, G. Munda, and R. Vreeker. Selecting an appropriate multi-criteria decision analysis technique for renewable energy planning. *Energy Sources, Part B*, 1(2):181–193, 2006.

[11] C. W. Potter, A. Archambault, and K. Westrick. Building a smarter smart grid through better renewable energy information. In *Power Systems Conference and Exposition*, pages 1–5. IEEE, 2009.

[12] A. Poullikkas. An overview of current and future sustainable gas turbine technologies. *Renewable and Sustainable Energy Reviews*, 9(5):409–443, 2005.

[13] L. Rao, X. Liu, L. Xie, and W. Liu. Minimizing electricity cost: Optimization of distributed internet data centers in a multi-electricity-market environment. In *INFOCOM*, pages 1–9, March 2010.

[14] X. Zhu, H. Chen, L. T. Yang, and S. Yin. Energy-aware rolling-horizon scheduling for real-time tasks in virtualized cloud data centers. In *High Performance Computing and Communications & IEEE International Conference on Embedded and Ubiquitous Computing (HPCC_EUC)*, pages 1119–1126. IEEE, 2013.