Trading power instead of energy in day-ahead electricity markets

Rens Philipsen[a,*], Germán Morales-España[b], Mathijs de Weerdt[a], Laurens de Vries[c]

[a] Faculty of Electrical Engineering, Mathematics, and Computer Science, TU Delft, the Netherlands
[b] Energy Transition Studies, Energy Research Centre of the Netherlands (ECN part of TNO), Amsterdam, the Netherlands
[c] Faculty of Technology, Policy, and Management, TU Delft, the Netherlands

HIGHLIGHTS

- Existing electricity markets, based on energy, are inefficient.
- We compare power-based scheduling to the current practice of increasing resolution.
- Power-based scheduling results in less reserve activation at equal computation time.
- Increasing the resolution requires great computational expense.
- We provide a set of market rules to implement and price power-based bids.

ARTICLE INFO

ABSTRACT

Day-ahead electricity markets are inefficient due to their coarse discretisation of time and their representation of electricity production and consumption in energy per time interval. This leads to excessive costs and infeasible schedules in the market clearing results. Some real-world systems have increased the resolution to improve accuracy, but this comes at a high computational cost. We propose an alternative, based on using linear power trajectories in the day-ahead scheduling process, which represent the momentary electricity production. Changing from a traditional energy-based to such a power-based formulation of the scheduling method can reduce cost by several percentage points, leading to a cost reduction of millions of euros in real-world systems on a yearly basis. Attempting to do so by increasing the resolution of the schedule would be accompanied by large increases in computational demands. Furthermore, we provide market design options to implement power-based bidding and pricing in day-ahead electricity markets, showing that pricing and market rules which encompass existing markets are readily available for implementation, and illustrate these with examples.

1. Introduction

Day-ahead scheduling in electricity systems is currently very inefficient. Irrespective of the precise way of scheduling (unit commitment (UC), one of its variants, or a self-dispatch market), the prevailing paradigm is that generators must meet the (largely inflexible) energy demand during each hour at minimal cost. To do so, the electricity system model must accurately incorporate the characteristics of the generators, the inflexible demand, and the varying output of renewable electricity sources. Only if these models are fully accurate we can make (day-ahead) schedules which ensure demand is met at minimal cost.

Current models, however, have some fundamental shortcomings. First, the discretisation of electricity demand to energy per hour (or other time interval) ignores information about momentary expected consumption, which may be available [1]. Since electricity markets are tasked with maintaining a momentary balance on the grid, ignoring such information in day-ahead markets might create inefficiencies, as it could lead to inefficient commitment decisions or infeasible schedules. Furthermore, this discretisation on the generators' side causes a misrepresentation of generator flexibility, and does not include any intra-hour flexibility which may be required to follow demand precisely [2]. As a result, the day-to-day operation of these systems is inefficient, increasing costs for end users.

These costs mainly arise from the need for more reserve capacity. Since day-ahead schedules are made on erroneous assumptions, overestimating generator flexibility and not representing momentary load, more reserve capacity must be on standby in real-time to compensate for deviations from the day-ahead schedule. We note that these problems exist in both centralised (such as in most North American systems) and self-scheduling systems (such as most European markets) [3].

---

* Corresponding author.
E-mail address: r.m.philipsen@tudelft.nl (R. Philipsen).

While in the former these inefficiencies lead to a less efficient dispatch, in the latter it leads to deterministic frequency deviations, which significantly impact system security and reliability [4]: a generation outage of 1300 MW causes a frequency dip of 50 mHz on the European synchronous zone [5], while market-related frequency dips (appearing on the turn of the hour, with smaller dips at the change of the programme time units at 15, 30, and 45 min) in the evening exceed 100 mHz *on average*. These temporary episodes of shortage and over-production are resolved via expensive reserves, intended to handle unforeseen events.

It is important to note that these problems are entirely due to the resulting electricity programmes not accounting for momentary production or consumption levels, but taking only energy per unit of time into account. These errors are therefore present in any system with electricity programmes which ignore momentary production, irrespective of how these electricity programmes are constructed: in day-ahead markets, via intra-day trading, or even through centralised scheduling in a fully deterministic scenario. Increasing or reducing the number of intermediary markets can therefore not solve this problem, nor should they: their task is to allow parties to deal with uncertainty.

A solution should ensure electricity programmes better represent momentary production and consumption. The current state of the art, which is already implemented in real-world systems by various means, is to increase the resolution of the electricity programmes, to the point where the energy programmes better represent momentary in- and output. This increase in resolution, however, comes at an exponentially increasing cost in computation. We, however, propose to redefine electricity programmes, explicitly basing them on momentary power in- and output.[1] In this paper, we compare power-based scheduling to increasing the resolution, and provide market design options to implement it in practice. Furthermore, we define a market which provides bid definitions and rules for bidding, market clearing, and pricing, all compatible with power-based scheduling. Thus, we deliver a set of rules for day-ahead markets which encourage momentary balance between supply and demand.

1.1. Contribution

The main contributions of this paper are as follows. First, we propose a market which is formulated in terms of linear power trajectories rather than energy per time unit. Using case studies based on five different real-world systems, with different sizes and fuel mixes, we quantify the potential efficiency gain of increasing the temporal resolution of scheduling or changing the scheduling paradigm from energy- to power-based. We analyse the impact of both increasing the resolution and changing to a power-based paradigm, analysing the trade-off between cost and system security and computational efficiency. We then discuss the barriers to implementing these solutions in a market environment, and propose a set of market rules which enable their implementation.

1.2. Organisation

First, we describe shortcomings in existing day-ahead scheduling methods. We select and describe two alternative solutions to these problems from literature, and discuss their respective pros and cons in Section 2. In Section 3, we perform a quantitative analysis of these alternatives based on various real-world systems. Section 4 describes the implementation of these solutions in a market environment. Section 5 draws the conclusions.

---

[1] Throughout this paper we mean by 'energy' and 'power' the physical quantities, as respectively measured in Joule and Joule per second. We do **not** mean to use either in the colloquial sense as a synonym for electricity.

2. Background

In this section we provide a brief review of shortcomings of day-ahead electricity markets, and review some previously suggested solutions.

2.1. Shortcomings in existing systems

Day-ahead scheduling in electricity systems is tasked with ensuring that sufficient generation capacity is available to meet demand at any given moment of the next day at minimal cost. To do so, time is discretised into programme time units (PTUs), and a balance between supply and demand is established for each PTU. This balance holds for the total amount of energy supplied and consumed during a PTU: we will therefore refer to it henceforth as energy-based scheduling. This balance does not, however, ensure that there is a momentary balance: as shown in Fig. 2, the momentary output of supply and demand may be entirely different all throughout the PTU, even if the energy balance holds for the total PTU.

Furthermore, the ramping capabilities of a generator must currently be described as a function of its total energy output during the previous PTU. However, there is no immediate link between the energy output during two consecutive PTUs; the ramping capability of a generator can only be described as a function of its output in power at a previous moment, not as a function of its output in energy during a previous period. As a result of no intra-PTU information being used, and physical constraints being misrepresented, day-ahead schedules are neither necessarily feasible nor efficient, even in systems with centralised real-time dispatch [1].

In systems with self dispatch, this leads to even bigger problems. In such systems, large deterministic frequency deviations can be observed at the turn of the PTU, as shown in Fig. 1. These deviations are caused by sharp changes in the ramping speed of generators at the turn of the PTU. As generators are incentivised (through imbalance penalisation) to adhere to their scheduled energy output only, they have no incentive to follow actual demand. Consequentially, they aim to fulfil their energy schedule as cheaply as possible, which in practice leads to momentary unbalance as inflexible demand does not make the same sharp changes, but follows a smooth trajectory instead. Any improvement to day-ahead scheduling which leads to better demand-following of generators can therefore also mitigate the side effects of self dispatch.

In real-world systems, a variety of methods are used in order to mitigate these problems. First of all, many systems employ a variety of markets after day-ahead schedules have been fixed. CAISO implements different ramping products to ensure the ramp requirements are met [6]. MISO has previously investigated longer dispatch intervals and additional ramping products [7]. Other solutions may include intra-day markets, of which various implementations exist in the European system [8], or real-time markets close to delivery. The core function of such markets, however, is to allow parties to adjust their position when imperfect information necessitates this. Furthermore, any inefficient commitment decisions can no longer be reversed on such short notice. Any changes to the market should therefore explicitly allow for enough time to start up or shut down generators where necessary.

2.2. Improving demand following in day-ahead markets

To effect better demand following, better representing demand and generation within the existing PTUs is paramount. There are three possible changes to day-ahead scheduling which can improve its accuracy: either we increase the resolution of the schedule, leading to more PTUs per day; we impose restrictions on the behaviour within a PTU after market clearing; or we change the market structure so that the market result better matches the physical reality.

Increasing the resolution of the schedule means a day is divided into more PTUs than is current practice. The length of each PTU therefore
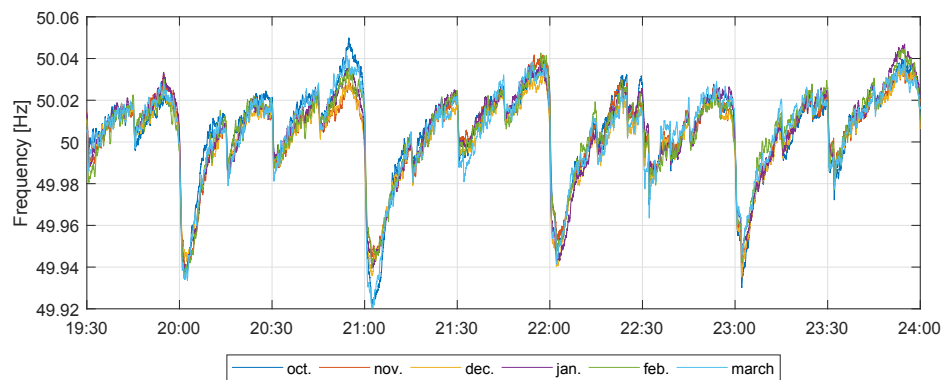
Fig. 1. Frequency on the European system 10–2016 to 03–2017, averaged over all days of the month. Based on data from [28].

decreases. The energy level in a smaller PTU better approximates momentary output, leading to shorter periods for which an energy balance must be guaranteed, better capturing the inflexibility of thermal units, and better assessing total costs [9]. This improvement in efficiency, however, would come at a computational cost [10]. Although this burden could be reduced, e.g. through decomposition methods [11], it is clear that increasing the number of decisions increases computational cost. An alternative approach is PTUs of variable sizes. This can be done through a rolling time horizon, using small PTUs for the near future while using larger PTUs further away, but these cannot be used to provide binding decisions beyond the smallest PTUs [12]. A more adaptive approach is difficult to implement in a market, as it requires prior knowledge of at the very least a demand curve [13]. The trade-off between cost efficiency and computational efficiency is therefore important to understand in order to judge this solution on its effectiveness; we examine this trade-off in Section 3.

A second alternative is to better approximate the intra-PTU output and consumption with more precise trajectories describing momentary output. These trajectories, if they correctly represent the technically feasible parameters of the unit, ensure the day-ahead schedules are in momentary balance. The most precise trajectories would describe the momentary power output. Matus et al. [14] proposes the use of Chebyshev polynomials to describe momentary power output. If these trajectories cover aggregated groups of small users, just as they are currently represented, these trajectories need not be very complex. The simplest approach to this is to impose linear intra-PTU trajectories. Such trajectories would mean all parties must have constant ramp rates within a PTU. ENTSO-E and Eurelectric [4] proposed to schedule based on energy, imposing linear intra-hour trajectories on the energy schedules. Rather than following demand precisely, this smoothens demand by imposing restrictions after market clearing. This has two disadvantages: imposing behaviour ex-post does not allow units to respond directly to market signals; and the imposed trajectories could deviate from the actual optimal solution as demand is not taken into account.

To avoid these shortcomings, scheduling methods could directly incorporate such trajectories [15]. This method is a bit stricter than the use of piece-wise intra-hour trajectories as proposed by Yang et al. [16], but clearly distinguishes resolution increases from methodological changes. The difference with existing scheduling methods is the enforcing of balance on the instantaneous power output (in MW) at the beginning and end of a PTU, and the imposed linear trajectory between these two points, rather enforcing this balance on the total output during the PTU (in MWh).

Fig. 2 shows the relation between energy and power schedules: all power profiles map to a single energy profile, but multiple power profiles may map to the same energy profile, as is illustrated with three distinct power schedules. As the area under the curves is equal, the same amount of energy is scheduled in all cases, but the schedule defines a precise trajectory for generators to follow. It is important to

notice that in this way the power-based scheduling approach not only satisfies a given energy demand, but also a given ramp demand which is not discernible in an energy schedule, while necessary to determine feasibility of a generator schedule. The examples of Section 4.3 also illustrate how ex-post changes to the market are not necessarily feasible, as the wrong generators may be committed day-ahead; for this reason we do not include this solution in our numerical evaluation in Section 3.

3. Case study through numerical simulation

In order to examine the trade-off between solution quality and computational cost, we analyse the two within-market solutions: increase in resolution, and changing to a power-based formulation. We compare these to the base case of energy-based scheduling with hourly resolution using a realistic case study of five different electricity systems. We use unit commitment (UC) models to represent these systems. We test whether the UC formulations are capable of providing the required electricity in real time by simulating real-time dispatch occurring every 5 min. These models allow us to measure the effect of the assumptions underlying the market models, rather than the difficult-to-model behaviour of parties within a self-dispatch market. By evaluating deterministic cases in real time, we test how well the UC models are able to deal with perfect information. We differentiate the runs as follows: by the model formulation, using EN for traditional energy-based planning and PW for power-based planning; by stage using UC for the unit commitment stage and RT for real-time dispatch; and by the resolution of the model, indicating the number of PTUs per clock-hour (1, 2, 3, 4, 6, or 12).

3.1. Experimental setup

We study five systems with different fuel mixes, both in terms of percentage of renewables as well as the composition of their fossil mix, resembling real-world systems. Fuel mixes are different for these systems to control for both long and short start-up trajectories. We ignore network constraints to isolate the effect on the different market designs, but in principle they can be added trivially. From Open Power System Data [17], we obtained time series with a 15-min resolution for Dutch electricity demand and renewable (both solar and wind) production, from which we constructed energy demand per five minutes through cubic interpolation with MatLab's *interp1* function, followed by least-squares linear interpolation to 5-min power intervals for the real-time dispatching stage. The generation portfolio (including capacity limits and fuel/generation type) is based on the DispaSET model [18] and a meta-study by DIW [19]. These are shown in Table 2. Furthermore, we made a few simplifying assumptions: Generator size is uniformly drawn from (100, 1000) MW for all fuel types; GHG-emission cost is €7.5/ ton; PJM was reduced in size; network constraints are ignored.
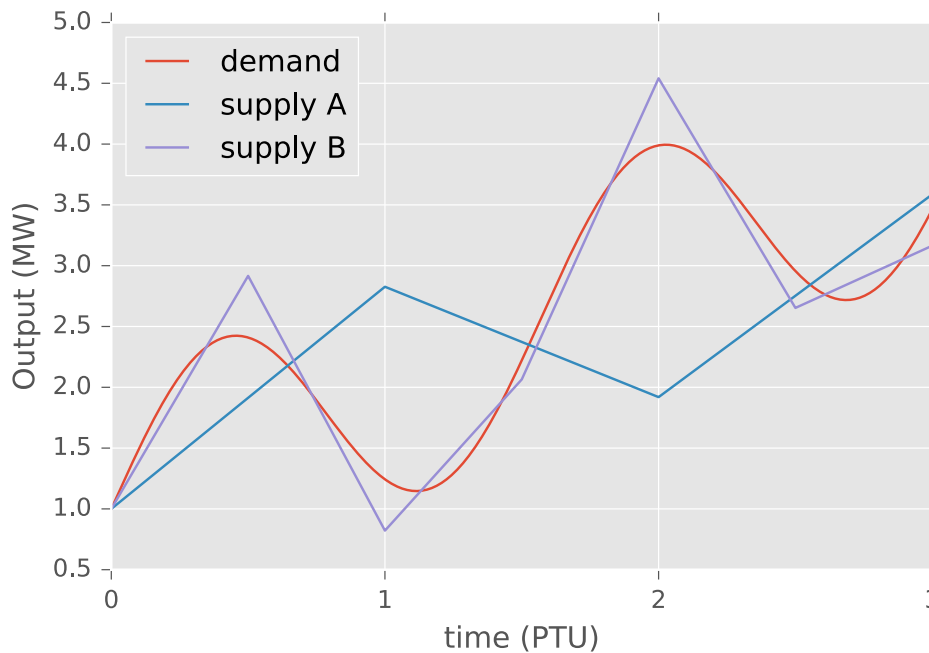
Fig. 2. Three power profiles with the same energy profile per PTU. Supply trajectory B also matches demand within each half-PTU.

3.1.1. Power vs Energy

We briefly recap the main differences between power-based and energy-based UC formulations. We use computationally efficient MIP models for both Energy [20] and Power [21]. These models include energy output during start-up and shutdown phases. No reserve constraints are imposed. In the energy-based models, the decision variables are the energy output of each generator during each time step and the renewable energy used, whose sum is constrained to be equal to the inflexible demand during the time step. In the power-based models, the decision variables are the power output of generators and renewable power at the end of a time step. Analogously, the supply–demand balance constraint is enforced at the end of the PTU. By the assumption of linear intra-period trajectories, a balance constraint at the beginning and end of any given PTU ensures momentary balance. Energy balance is therefore implied by these power balance constraints. For both models, renewable energy can be curtailed (at a penalty of 50€/MWh) in both the UC and RT stage. Furthermore, energy can remain unserved in RT if necessary to ensure feasibility, but is penalised at 10000€/MWh. Note that this is not possible in the UC stage: any unserved energy is therefore entirely due to the UC stage being unable to account for known events.

3.1.2. Model runs

We compare the performance of these two formulations by making a day-ahead schedule at a given number of UC intervals per hour, and then simulating the real-time execution of this schedule with a finer resolution of 5 min. We vary the number of UC intervals per hour from 1 interval/hour (which matches UC with an hourly resolution) to 12 intervals/hour (which matches setting commitment decisions every 5 min, at the same time scale as the real-time dispatch). By using the

same demand profile in both stages in a deterministic setting, we are able to assess which formulation better deals with perfectly known events.

The UC models are solved for 48 h to prevent end-of-horizon effects from influencing the results: the commitment decisions for the first 24 h are evaluated. Initial conditions are derived from a lead-in day (starting with all generators turned off) for the first day, and for subsequent days the system state after 24 h of the previous day's UC are used. The model runs are fully deterministic: there is no uncertainty about demand or renewable supply. Consequentially, we do not include the scheduling of reserves in the formulation.

3.1.3. Scenarios

We consider a number of different systems, each representing a real-world system. Each system has a different level of RES penetration, and a different fossil fuel mix (leading to more/fewer highly flexible plants), as shown in Table 1. The generators for these systems are generated pseudo-randomly, where the maximum capacity and variable cost have an uncertain range of 10% around the listed figures. The systems based on Danish and German (DK, DE) portfolios have a relatively high RES penetration, while we can classify the Dutch and British (NL, GB) systems as medium-RES. PJM has little installed RES capacity. The demand and RES time series are taken from Dutch data of 2015 [17], and scaled according to full system size and installed capacity. We evaluate the models based on the data of two weeks: the week with the highest ramping requirement of the dataset, and the week with the highest variance in residual demand.

3.1.4. Performance indicators

We evaluate the formulations primarily by system security. We use

Table 1
System fuel mixes in percentages.

System	Total [MW]	Fossil [MW]	# generators	Lignite	Hard coal	Oil	Natural gas	Nuclear	Wind	Solar	Bio
NL	34,470	29937.2	58	0	23.93	0	58.34	1.42	9.14	3.89	1.19
DK	11,550	5999.1	16	0	20.87	0.43	18.35	0.09	42.6	5.37	12.29
GB	80,946	65743.53	127	0	25.58	3.56	39.43	11.62	13.5	0	1.03
DE	195,780	105573.9	192	10.85	14.73	1.37	14.9	6.17	19.92	20.67	3.11
PJM	17162.8	16239.6	39	0	36.21	6.51	31.22	19.71	0.51	0.09	0.97

Table 2
Plant characteristics (based on [18,19]).

Fuel	Min output [%P_{max}]	Ramp Up [MW]	Ramp Down [MW]	Min Up Time [h]	Min Down Time [h]	SU Cost [€]	SU Time [h]	Efficiency [%]	GHG [ton/ MWh]	Fuel cost [€/ MWh$_{in}$]
Coal	30	300	300	8	6	11724	7	0.395	0,95	9.3
Gas	40	670	670	5	1	11254	1	0.525	0,45	38.3
Oil	30	300	300	8	6	3528	7	0.41	0,79	20
Nuclear	90	206	206	24	24	20000	12	0.405	0	37
Lignite	30	300	300	8	6	13685	7	0.385	1	20
Bio	30	300	300	8	6	2000	7	0.45	0	25.8

the curtailed load during real-time dispatch as a measure of how well the unit commitment formulation is able to deal with perfect certainty. in real systems, reserves would of course be in place to handle such events, and load would not necessarily be curtailed. As setting reserve capacity limits is a somewhat arbitrary decision and would impact the efficiency of the scheduled generators, and there is no uncertainty to be taken into account, reserves form no part of our UC formulations. We could very well interpret load curtailment figures as reserve activation; it is that part of demand which was met by the available generators in the UC formulation, but not in the RT evaluation of that UC schedule. It is a measure of how well the UC formulations are able to deal with certainty (as there is no stochasticity involved in our experiments at all).

Total cost is largely dominated by the arbitrarily imposed cost of this curtailed load, and we therefore only consider cost as a metric for those cases where there is no load curtailment. Finally, we are also interested in the computational efficiency, as indicated by the total run time.

Two important benchmarks are EN with resolution 1, representing current practice in real-world power systems, and PW with resolution 12, where day-ahead schedules are made using the real-time dispatch method, and is therefore the best schedule possible. We therefore use the former as a benchmark for runtime, and the latter for total cost and curtailed load.

3.2. Results

3.2.1. Curtailed load

Figs. 3–6 show the performance of real-time execution relative to the optimal benchmark of PW-12, differentiated by the different UC formulations and the resolution of the UC stage. Each data point indicates the performance for a single day. Less is of course better, as ensuring security of supply is an important goal of electricity markets. For DK and DE, the high-RES systems, the difference is the clearest.
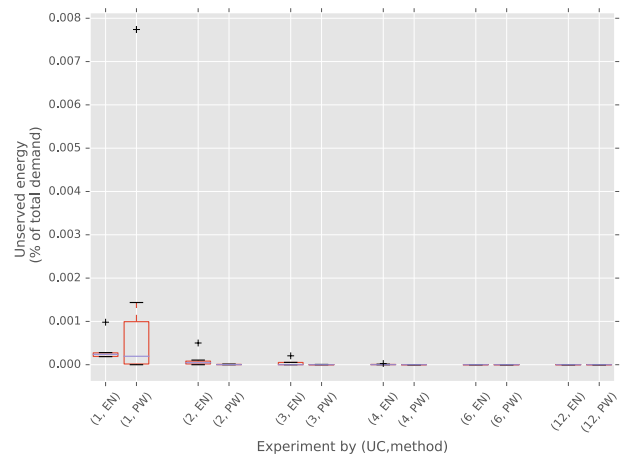
Clearly the PW outperforms the EN at each resolution, and the trend
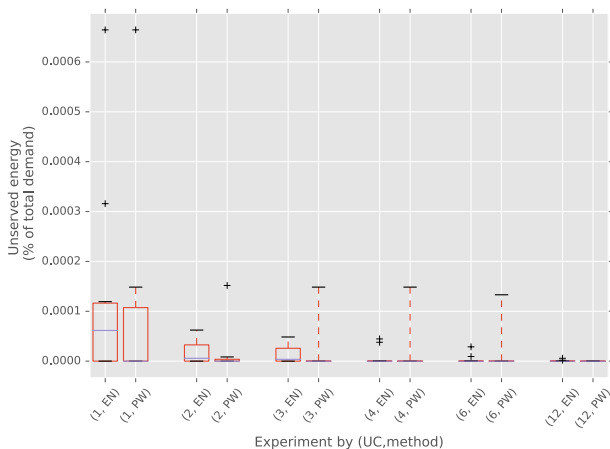
Fig. 3. Unserved energy in the DK system.
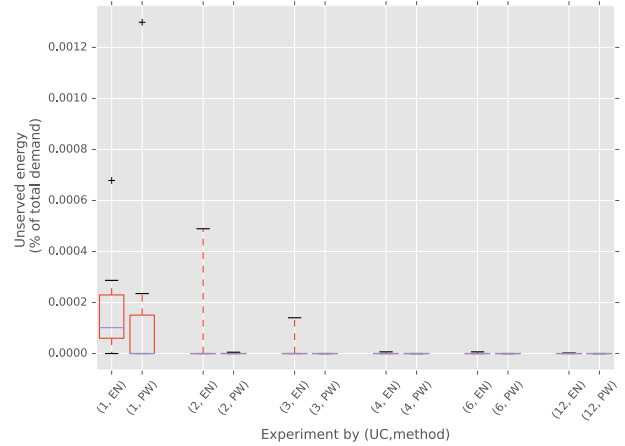
Fig. 4. Unserved energy in the NL system.

Fig. 5. Unserved energy in the GB system.

of lower costs at higher resolutions is also sensible, as higher UC resolutions better approximate the real-time execution stage. At a resolution of 12 PTUs per hour, the PTU length is 5 min and therefore exactly matches the real-time execution phase. It is noteworthy that even at this stage, the EN-formulation is unable to make minimal-cost schedules which avoid load curtailment (as in the DK and DE systems).

In the high-RES scenarios, there is a clear improvement of PW over EN. In the DK system, PW outperforms EN across all resolutions, while for DE there is a single outlier at PW-1 (but the overall trend is still a decrease of curtailed load as resolution increases, and of PW outperforming EN). For systems NL and GB, the medium-RES systems, PW quickly reduces its unserved energy to zero, while the EN formulation must shed load even at higher resolutions. The results are mostly size-invariant: GB is essentially a larger NL, and DE a larger DK. The difference between different levels of RES capacity, however, are clear.
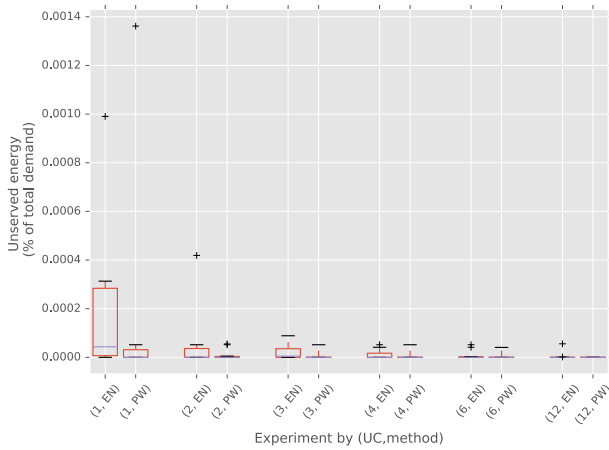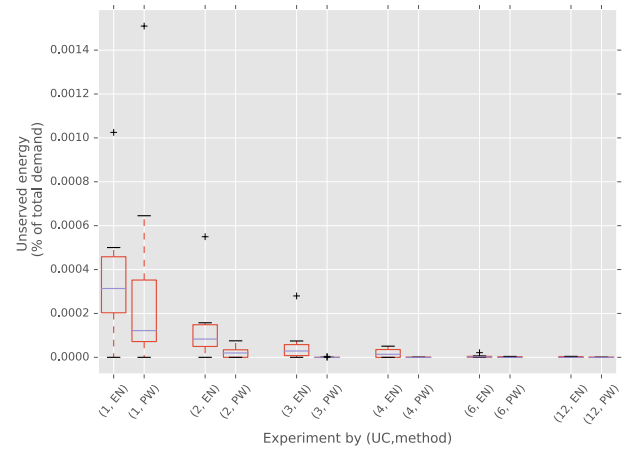
Fig. 6. Unserved energy in the PJM system.

Fig. 7. Unserved energy in the DE system.

3.2.2. Runtime

Runtime, meanwhile, is mostly invariant to the formulation used, but greatly increases as resolution increases, as is summarised in Fig. 9 for all systems on a logarithmic scale. At first, computational cost increases slowly as resolution increases, but the run time for EN-12 and PW-12 reaches a thousandfold increase over EN-1. This is a rather straightforward result: computational time increases quickly due to the increasing number of integer variables of the MIP formulation. What is of more practical interest, however, is to examine the combination of run time and curtailed load. For DK, it shows that the PW-2 formulation delivers results which are already very close to the PW-12 solution, and outperforms EN-6. EN-6 also has a computational cost increase of about 100x. If we keep in mind that currently Euphemia has a time limit of about 40 min [22], it is clear that such large increases in run time are not acceptable. Given computational limits, these results show that the only way to improve system security is to change the formulation from energy-based to power-based.

3.2.3. Total cost

As the load curtailment cost component largely dominates the total cost, we only assess the total cost of those cases where no load is curtailed. These runs are shown in Fig. 8, displaying the relative performance in terms of total cost (the sum of variable cost of generation, startup costs, and wind curtailment). Overall, PW outperforms EN slightly, although this not the case in a pairwise comparison. A close examination of the experimental results suggests this is due to the fact that startup decisions may fall on different days, depending on the formulation. For example, if EN starts a generator on day 2, and PW does not start that same generator until day 3, EN will likely outperform PW on day 3, simply because the startup costs have already been accounted for on day 2.

The qualitative conclusion is therefore that EN does not outperform PW in cost efficiency. Quantitatively, PW has a slight advantage. Although the advantage is certainly minor in relative terms, it must be pointed out that the total electricity consumption of Germany was approximately 600 TWh in 2016, and a relative performance increase of even 0.1% in absolute costs is in the order of magnitude of €20 million.

Overall, we see power-based scheduling has a number of advantages. It outperforms energy-based scheduling in terms of both remaining imbalance (Figs. 3–7) and cost (Fig. 8). An equally important observation is that it delivers these advantages without an increase in computational cost, as shown in Fig. 9. These results make power-based scheduling a superior scheduling method for electricity markets.

Few real-world systems, however, are fully centralised. In most systems, markets are employed to match supply and demand. In the next section, we therefore explore which changes must be made to the
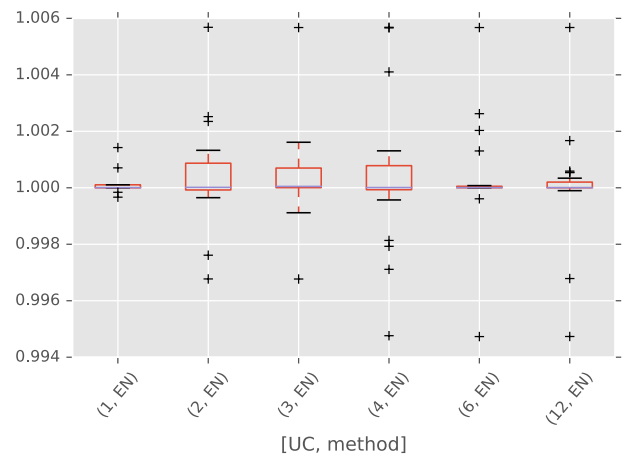
Fig. 8. Total cost (EN divided by PW) in experiments without curtailed load, by UC resolution.
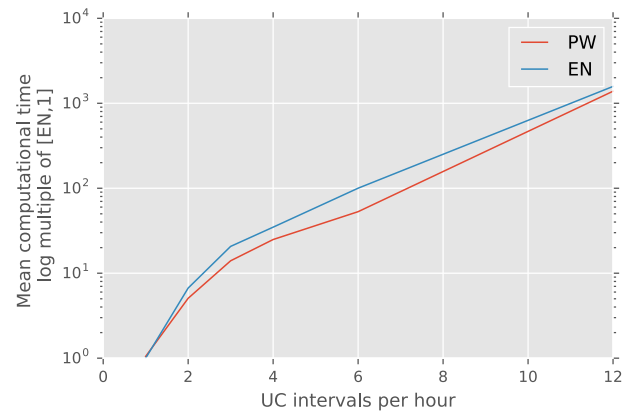
Fig. 9. Average UC runtime increase over EN-1 on logarithmic scale for all systems.

market design in order to successfully implement power-based scheduling in practice.

4. Market design

The implementation of a higher resolution of the day-ahead market is rather simple: no fundamental changes are necessary, it suffices to increase the resolution of the market clearing algorithm. The

implementation of power-based scheduling into a market, on the other hand, is more complicated. Existing markets communicate energy profiles to the market participants, who are expected to adhere to these energy programmes. In a power-based market, however, market operators should communicate power trajectories rather than energy profiles. This requirement, in turn, may affect other aspects of the market design as well. In this section, we discuss the changes necessary for the successful implementation of power-based scheduling in a market environment, including bidding rules, clearing rules, and pricing rules.

In a centralised scheduling and dispatch system, implementation would be relatively straightforward. It would suffice to have both generators and demand communicate all their characteristics to the market operator. This is existing practice in some systems (e.g. US systems), but we realise this may not be desirable or feasible in others. European systems, for example, work based on the Euphemia algorithm, which is used for day-ahead market clearing in Europe [23]. We therefore turn our attention to how these improvements to day-ahead scheduling could be implemented if we require that market parties bid and trade with each other, as is current practice.

4.1. Market rules for power-based scheduling

We start by describing existing market rules. We will then show that the market clearing result is insufficient, and discuss required changes. Changes to the clearing rule, in turn, require changes to the definition of bids and the pricing rule, where for both we aim to stay as close to the existing market (in terms of properties of the pricing rule and freedom for the market parties) as possible. We define multi-period bids as our starting point for bids, where parties communicate their physical constraints (such as ramping or output limits) and variable costs. The market is cleared for a fixed horizon (i.e. the next day), and communicates the precise power trajectories all parties should follow. We enumerate the criteria for a pricing rule based on existing markets, and define prices as a price per megawatt of output at the end of each PTU. We then provide some examples of how these prices might seem counterintuitive, but still guarantee adherence to these criteria.

Our foremost goal is to show that, given the characteristics of existing markets (which we aim to improve upon), it is possible to design a consistent set of market rules which implement power-based scheduling. In doing so, we aim to describe the minimal requirements to which bids and bidding rules must conform. Further improvements may therefore well be possible.

4.1.1. Existing market design

European day-ahead electricity markets currently operate as follows. The day-ahead market is cleared 12–36 h in advance of delivery. Bids are submitted through the national electricity market operators (NEMOs) to a single European market clearing algorithm, Euphemia [23]. The core of the algorithm is a constrained welfare maximisation problem. As the submitted bids are inputs for the objective function and constraints of the optimisation problem, Euphemia defines what permissible bids consist of. The atomic bid is a limit order for an amount of energy [MWh] during a specified hour [h] at a specified price [€/MWh]. More complex bids can extend or combine such bids by creating dependencies or conditional acceptance; we review these possibilities under *bidding rules*. The communicated result is the amount of energy each party is supposed to deliver or consume during each hour of the day; and a price for energy during each hour, which is the marginal price of energy delivered during that hour. The concept of marginal cost pricing is well-established in the power systems community. These so-called shadow prices, obtained from the dual variables of optimisation constraints in linear programming, form a competitive equilibrium and ensure envy-free prices for all participants, who all recover their costs [24].

4.1.2. Clearing rule

The market result of existing markets, however, does not provide generators with sufficient information to follow a power trajectory; the clearing result only communicates an amount of energy. This is therefore our first criterion for a new market design: the market result provides each participant with a piecewise-linear power trajectory to follow.

4.1.3. Bid definition

In order to provide a power trajectory to bidders, the market operator must receive precise information on the physical capabilities of the bidders. Bids allow both generators and demand to provide these required inputs to the scheduling problem. This means they must be able to express their physical characteristics with enough precision to allow for power trajectories as an output. In order to reach the desired outcome, existing bids no longer suffice; the linear ramping constraint is not included in existing bid definitions. Furthermore, due to the requirement of linear ramping during a PTU, bids for a single PTU no longer exist. Consequentially, we require a bid which explicitly links two subsequent PTUs, and incorporates the linear ramping constraint.

A bid must at least consist of a lower and upper production limit at the start/end of a PTU [MW]; a maximal ramp rate, both up and downwards [MW/h]; and a price per unit of energy [€/MWh].

Using this bid definition, PTUs are linked to each other through the ramp rates, which set limits on the intra-hour gradient of the resulting power trajectory. Production limits, as in existing markets, are necessary to ensure feasibility.

4.1.4. Pricing rule

In existing markets, the price of energy is the marginal price of energy during that hour. We briefly recap the desired characteristics to which a pricing rule must (and currently does) conform. For the pricing rule, we aim for the following characteristics. These characteristics are well-established in micro-economic theory.

- Efficient: Given all bids, social welfare is maximised.
- Incentive compatible (or truthful): honestly reporting your preferences is the best strategy (i.e. it maximises your own utility) given that others do so as well (known as Bayes-Nash incentive compatible) or irrespective of what others do (dominant-strategy incentive compatible).
- (Ex-post) Individually rational: mechanism participants receive non-negative utility from participating, in all possible states of the market, implying costs are always recovered for all bidders.
- Budget balanced: a mechanism is strongly budget-balanced if the sum of payments by the market operator is precisely 0.
- Envy-free: No generator values another production profile and corresponding payment higher than the profile and payment it is assigned.

It is impossible, however, to satisfy all five. Due to the impossibility theorem by Myerson and Satterthwaite [25], the first four can never be satisfied by a pricing mechanism simultaneously. Although it is possible to design a mechanism that addresses incentive compatibility, such as the celebrated Vickrey-Clarke-Groves mechanism [26], such a mechanism will have different downsides due to the impossibility theorems of Myerson and Satterthwaite [25] and Green and Laffont [27]. In practice, applying VCG would lead to discriminatory pricing, benefiting larger generating companies (as each participant is paid according to their impact on the total welfare of all other participants, and the impact of a large company is bigger than that of a smaller one), and might run a deficit. Existing markets face the same problem, and we therefore follow their example of relaxing the criterion of incentive compatibility. Implementing marginal pricing would ensure the four remaining criteria are met, as they are fundamental properties of marginal pricing in any situation.

In implementing marginal pricing, however, we must overcome one serious hurdle. In an energy-based market, the balance constraints are enforced on the energy output, and the marginal cost of supplying an additional unit of energy during a PTU can be expressed in €/MWh, resulting in a price for that PTU alone. In a power-based market, however, PTUs are irrevocably linked to each other due to the linear power trajectories, where increasing the output of a generator at the end of one PTU increases the amount of energy in both the preceding and in the subsequent PTU. Although these trajectories correspond to a unique energy profile, we can no longer price the energy profiles since they are no longer independent. Instead of pricing the energy output during the PTU (in MWh), we must price the power output at the end of the PTU (in MW). This uses the dual variable of the power balance constraint at the end of a PTU, instead of the dual variable of the energy balance constraint during the PTU, as the price for power delivery at that moment. By doing so we can maintain the properties described above, but the price is now in €/MW.

Such prices may lead to a seemingly paradoxical situation. If each bidder now pays their output at the end of the PTU times the price, the situation arises where a generator which shuts down will receive no payment for energy supplied during its last operational hour. It should be noted, however, that when ramping upwards, generators are paid more than their electricity output warrants. Since a generator which ramps down must have ramped upwards at some earlier moment, variable costs have already been recovered by the payments in the previous hours. The apparent paradox is therefore illusory, as we will further illustrate with a numerical example in the next section: marginal pricing of power profiles is fully individually rational.

4.1.5. Bidding rules

In designing bidding rules, we aim to describe the minimal set of rules which ensure that bidders have at least the flexibility they currently have. Currently, Euphemia accepts (and by extension, NEMOs offer) a number of bid products which extend the previously-described atomic bid. These bids, linking PTUs, are known as complex bids. We briefly list them, describe them, and discuss their incorporation in a power-based market.

Load gradient orders. Load gradient orders describe dependencies between PTUs to capture multi-period ramping constraints. They consist of bids for sequential PTUs, linked by ramping constraints. Clearly, this is captured by our description of a single bid. Block orders describe subsequent PTUs for which a generator either runs during all PTUs, or during none. Using our bid definitions, this can also be accomplished through the minimal output levels of individual bids. Furthermore, by bidding the appropriate ramping limits, participants can ensure they run for multiple hours if their bid is accepted.

Flexible hourly orders. A flexible hourly order is an order for an amount of energy delivered during a single PTU, where the precise PTU is to be determined by the market outcome. This is the only bid type we cannot cover in power-based scheduling, which is by design: as we impose linear ramping during a PTU, it is no longer possible to start and end a PTU at zero and still consume electricity. The alternatives for a party desiring to place such a bid would be to replicate it for two subsequent PTUs, or to turn to bilateral trading to ensure the desired power profile can be followed.

Minimum income conditions. Minimum income conditions (MIC) are bids which require fixed (start-up) costs to be covered. They impose restrictions on how multiple subsequent sub-bids are accepted or rejected. There are no barriers to imposing such restrictions on our power-base bids.

Linked orders. Finally, linked orders can be introduced by including conditional acceptance between bids. Such bids can then be used to

capture price differences between operating ranges, but their use is likely to be lower than in current markets due to the fact that our basic bid already allows for restrictions on the range within which a plant can be operated.

4.2. Implementation issues

So far, we have discussed the changes in market rules necessary to implement power-based scheduling in a day-ahead market. There are a two issues which fall outside the scope of this work, but which we still touch upon briefly. First we discuss the impact a new market paradigm has on flexible demand, within the context of the day-ahead market (therefore ignoring its role in other ancillary markets); secondly, we discuss imbalance settlement.

4.2.1. Role of hydro, storage and flexible demand

In essence, there is little difference between hydro-power and flexible demand: both, to a certain physical limit, allow for intertemporal shifting; tend to have high flexibility (i.e. high ramp rates); and by this virtue, can increase the efficiency of electricity schedules if incentivised properly. Under an energy-based market, there would be no incentive to follow actual demand, and self-scheduling of hydro-power plants and flexible demand could exacerbate momentary imbalances by quickly ramping up in order to meet their full energy profile [4]. Creating incentives to follow linear power profiles reduces the negative impact they have on the system, and are able to contribute to further system stability by providing reserve capacity to compensate for actual imbalances.

Systems with highly frequent (e.g. every 5 min, as in the US) real-time centralised dispatch, of course, face no such problems. For these systems, however, there are potential efficiency gains because flexibility now has a day-ahead value. These efficiency gains are also directly shown in our numerical results, where we dispatched all units every five minutes in a central and optimal manner (simulating US real-time markets). Using such dispatch, there are no deterministic imbalances or frequency deviations, but still we showed efficiency gains.

Since both hydro-storage and flexible demand are constrained in their total production and consumption (due to reservoir limits, or total consumption limits), there are intertemporal dependencies: using a flexible resource during one PTU means it is no longer available during the next. In order to efficiently plan when these sources of flexibility should be available, it is necessary to have a good understanding of the flexibility of all other parties. Power-based scheduling, more so than energy-based scheduling, provides this insight, irrespective of the dispatch method.

4.2.2. Imbalance settlement

Once the day-ahead schedules are known, how should deviations be dealt with? This question, which we have barely touched upon in this paper, is essential for the proper implementation of power-based scheduling into a market. A well-functioning imbalance settlement mechanism is imperative to incentivise market parties to actually follow the power profile they have been assigned. Clearly, the existing imbalance settlement mechanism of measuring the total amount of energy during a PTU no longer suffices, but an appropriate imbalance mechanism is yet to be designed.

4.3. Examples

We illustrate the application of marginal pricing to power output at the end of a PTU with some examples. We consider two scenarios, SC1 and SC2, in a system with three generating units, whose data is shown in Table 3, and the demand presented in Table 4 and Figs. 10 and 11. Note that in the rows indicating demand in power, the values represent the power demand *at the end of the hour*, while in the rows showing demand in energy, they represent the energy demand *during the hour*, as

Table 3
Generator data.

Name	Min [MW]	Max [MW]	Ramp rate [MW/h]	Offer price [€/MWh]
G1	0	200	65	25
G2	0	100	120	37
G3	0	60	180	90

Table 4
Demand data. Power in MW at the end of the hour, energy in MWh during the hour.

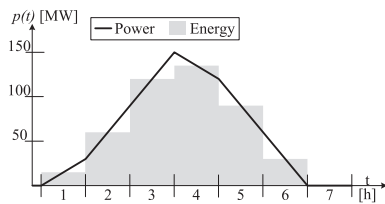| Hour	0	1	2	3	4	5	6	7
Power SC1	0	30	90	150	120	60	0	0
Power SC2	0	0	120	120	150	30	30	0
Energy SC1	–	15	60	120	135	90	30	0
Energy SC2	–	0	60	120	135	90	30	15

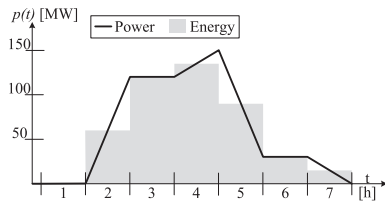
Fig. 10. Demand in Scenario 1.

Fig. 11. Demand in Scenario 2.

is also shown in Figs. 10 and 11. This example consists of two conventional generators (G1 and G2), and one fast-start unit (G3) which provides system reserves if needed in real-time operation. We present two examples: in SC1, we show that with power-based pricing costs are recovered by the generators, even if the power output drops to zero at the end of some hour during which they were producing. In SC2, we show that the power-based market reaches a feasible solution where the energy-based market does not, and as a result the power-based market leads to a more efficient result in actual dispatch.

4.3.1. Scenario 1

Clearing the power-based market in SC1, G1 can completely supply the required power trajectory, as shown in Table 5, which also lists the associated energy profile, and the marginal prices in both markets. This

Table 5
SC1: Dispatch and prices for both power- and energy-based markets.

Hour	0	1	2	3	4	5	6	7
G1 [MW]	0	30	90	150	120	60	0	0
G2 [MW]	0	0	0	0	0	0	0	0
G3 fast-start [MW]	0	0	0	0	0	0	0	0
G1 [MWh]	0	15	60	120	135	90	30	0
Energy-based prices [€/MWh]	0	25	25	25	25	25	25	0
Power-based prices [€/MW]	0	25	25	25	25	25	0	0

Table 6
SC1: Costs and income in the day-ahead market.

	Energy-based		Power-based	
	Costs [€]	Income [€]	Costs [€]	Income [€]
G1	11,250	11,250	11,250	11,250
G2	0	0	0	0
G3 fast-start	0	0	0	0

is also the cheapest solution. For all hours at the end of which the amount of power traded is larger than 0, the price is equal to the cost of the marginal unit, 25€/MW. Clearing an energy-based market, G1 is also able to provide the entire demand, although it is not specified which power trajectory the unit should follow. Since G1 is the only unit providing the demand, the marginal energy price for all hours during which energy is traded is again 25€/MWh.

Upon first inspection of the power-based market result, G1 does not seem to recover its operational cost, as there are fewer hours during which it earns any income. G1 does recover its operational costs as is shown in Table 6.[2] This is due to the fact that the costs are recovered in the earlier hours, where the power supplied at the end of an hour is larger than the energy supplied during the hour because the generator is ramping up, as could be seen in Table 5.

4.3.2. Scenario 2

In an energy-based market, G1 is expected to easily follow the demand profile, as all energy ramps are below 65, and the energy demand is the same in hours 2–6 as in SC1. G1 is therefore dispatched as in scenario 1 in hours 2–6, and consequently we see the same price result. This day-ahead schedule, however, is infeasible.

If we attempt to execute the day-ahead schedule, G1 cannot provide the energy ramp of 60 MWh in hour 2 because the generator was not already ramping up, and the actual ramp requirement was 120 MW (see Fig. 11). To supply demand we must rely on reserves provided by G3, as shown in Table 7. G3, of course, is more expensive than G2, which could have provided the required production profile had it been committed in the day-ahead market. Furthermore, we see that the day-ahead market did not correctly account for the need of ramp-down capacity in hour 5. Since G1 cannot ramp down fast enough to completely follow demand, it is necessary to use G3 to provide reserves again in hour 4 and 5. This optimal redispatch is shown in Fig. 12.

The power-based market correctly represents these physical limitations, and schedules as shown in Table 8: it commits both G1 and G2, scheduling G2 where necessary to maintain an adequate ramping speed. Fig. 12 shows the optimal dispatch, with the exception that G2 now provides the power instead of G3. The marginal prices (shown in Table 8) also deserve a closer look. At the end of hours 2 and 4, G2 is the marginal generator, and therefore it sets the price. What stands out, however, is the low price at the end of hour 5. Due to ramping constraints, an additional unit of power provided at that moment would allow us to reduce the output of G2 by 1 at the end of hour 4, as there is less ramp-down demanded from G1. Consequently, G1 provides additional power at the end of both hours 4 and 5, delivering in total 2 additional MWh of electricity during hours 4–6. The total electricity output of G2 is reduced by 1 MWh. The marginal cost of this 1 MW increased demand is therefore $2 \cdot 25 - 37 = 13$. Even though this price is below the electricity production cost of all generators, costs are still completely recovered, as is shown in Table 9. Moreover, the cost of dispatch matches the costs in the power-based day-ahead market, which was not the case for the energy-based market.

---

[2] The income is simply the sum-product of the dispatch rows and the corresponding prices in Table 5.

Table 7
SC2: Actual dispatch [MW] for energy-based market, and day-ahead prices.

	0	1	2	3	4	5	6	7
G1	0	0	65	120	95	30	30	0
G2	0	0	0	0	0	0	0	0
G3	0	0	55	0	55	0	0	0
Price [€/MWh]	0	0	25	25	25	25	25	25
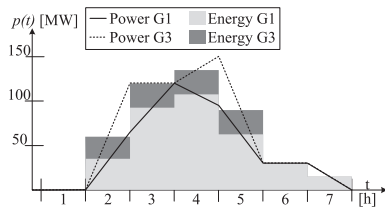
Fig. 12. Actual dispatch in Scenario 2.

Table 8
SC2: Power-based schedule [MW] and prices.

Hour	0	1	2	3	4	5	6	7
G1	0	0	65	120	95	30	30	0
G2	0	0	55	0	55	0	0	0
G3	0	0	0	0	0	0	0	0
Price [€/MW]	0	0	37	25	37	13	25	0

Table 9
SC2: Costs and money recovered according to day ahead, and the cost of actual dispatch.

	Energy-based			Power-based		
	Day-ahead cost [€]	Day-ahead income [€]	Dispatch Cost [€]	Day-ahead cost [€]	Day-ahead income [€]	Dispatch Cost [€]
G1	11,250	11,250	8500	8500	10,060	8500
G2	0	0	0	4070	4070	4070
G3	0	0	9900	0	0	0
Total	11,250	11,250	18,400	12,570	14,130	12,570

5. Conclusion

Existing electricity markets are inefficient: their very design makes it so. Markets let parties trade energy, to be delivered anywhere within a programme time unit. Ensuring energy balance during such a time unit, unfortunately, does not guarantee the momentary balance of supply and demand which is a necessary condition for the safe and reliable operation of electrical power systems. Energy-based scheduling methods consistently misrepresent the flexibility of generators. Since average energy levels do not provide any information about the momentary output of a generator, and ramping limits are dependent on momentary output rather than average levels, traditional energy-based scheduling overestimates their flexibility. In practice, this means generators must misrepresent their ramp rate when a market is cleared simply to ensure the resulting schedule is feasible. Energy-based markets, in short, trade the wrong product.

Scheduling electricity in terms of power, rather than energy, can improve electricity markets in a number of ways. Primarily, it improves the efficiency of day-ahead schedules at a lower computational cost than increasing the resolution does. Day-ahead schedules based on power trajectories are, in all ways, superior to energy-based schedules.

The reduction in load curtailment (or reserve activation) runs in the order of a few percentage points. This reduces costs for consumers, and reduces the need for online thermal power plants, improving the competitive position of renewable electricity sources. In self-dispatch systems, power-based scheduling also alleviates the imbalances which follow from the existing market design, preventing frequency shocks which threaten security of supply and freeing up expensive reserves, and resulting in more efficient schedules in general. By committing to a specific trajectory (i.e. always linear), efficiency might be lost in some very specific cases, such as when a cheap generator might be able to reach its maximal output during a programme time unit but this would result in a non-linear trajectory. The gains, arising from better representation of both the momentary inflexible demand and generation, as well as a better (if overly conservative) representation of the actual constraints on generators outweigh the losses, as the experimental results corroborate. The experimental results furthermore indicate that a power-based market's relative performance increases when the share of renewable electricity sources in the system increases, which is an important conclusion with an eye on the policy goals of most industrialised countries.

Coordinating power-based trajectories in a market was until now an unresolved problem. We have defined a market which fills this gap by providing bid definitions and rules for bidding, market clearing, and pricing, delivering a set of rules for day-ahead markets which encourage momentary balance between supply and demand. The proposed market model is very much in line with the operation of existing markets, making only the minimal changes necessary to fully capture the advantages of power-based scheduling. In doing so, it improves economic efficiency and makes way for further integration of RES.

In our experiments, we assumed the absence of transmission constraints in order to show as clearly as possible that the inefficiency of schedules is solely due to the assumptions underlying energy-based schedules. This can easily be extended to include transmission constraints. Transmission constraints are also more accurately expressed in power than in energy. When transmission constraints are expressed in terms of energy rather than power, transmission lines face the same problem as a generator reaching its maximal output: a line rated at 100 MWh per hour may require flows in excess of 100 MW to actually transport 100 MWh. Therefore, power-based network constraints can also improve transmission congestion management. Power-based scheduling incorporates the true transport capacity, reducing the margin of error the system operator must take into account, therewith increasing the efficiency of the system. Our proposed power-based market design can easily be extended to include transmission constraints. Power-based locational marginal prices (LMPs) can be used to include network constraints. Power-based LMPs are obtained in the same (mathematical) way as the LMPs in energy-based markets. The difference is the meaning of these power-based LMPs, which is exactly the same meaning of the prices we present in this paper: they reflect the cost of providing 1 MW of power at a given location at the end of a programme time unit or period.

Looking ahead in this line of research, the market design requires rethinking of the imbalance penalisation rules. Existing imbalance settlement in self-scheduling markets does not incentivise generators to actually follow the assigned power trajectory. These imbalance rules must therefore be redesigned to account for momentary power rather than total energy per time unit.

Furthermore, market manipulation is important to investigate. Inefficiencies in day-ahead scheduling create a need for reserves and balancing. At least in theory, it is possible in a self-dispatch system to create momentary imbalances, for which the system operator will then reward whoever resolves it. In a power-based market, the market result is a momentary schedule from which, in theory, it is more difficult to profitably deviate. The manipulation of such a market will of course strongly depend on its relation with other markets, including the imbalance settlement mechanism referred to previously.

Appendix A. UC Formulations

In this Appendix, we briefly describe the unit commitment formulations used in the experiments. First we introduce the notation used in the UC formulations.

Indices and Sets

$g \in \mathcal{G}$	Generating units, running from 1 to G
$\mathcal{G}^{S} \subseteq \mathcal{G}$	Slow-start generating units
$\mathcal{G}^{1} \subseteq \mathcal{G}$	Generators with $TU_g = 1$
$s \in \mathcal{S}_g$	Startup segments, running from 1 (the hottest) to S_g (the coldest)
$t \in \mathcal{T}$	Hourly periods, running from 1 to T hours
$i \in [1, 2, ...]$	intervals of startup and shutdown trajectories.

Parameters

C_g^{LV}	Linear variable production cost [€/MWh]
C_g^{NL}	No-load cost [€/h]
C_g^{SD}	Shutdown cost [€]
C_{gst}^{SU}	Startup cost for segment s [€]
C^{VW}	Variable production cost (bid) of wind [€/MWh]
D_t^{E}	Energy demand for hour t [MWh]
D_t^{P}	Power demand at the end of hour t [MW]
$\overline{P}_g$	Maximum power output [MW]
$\underline{P}_g$	Minimum power output [MW]
E_{gi}^{SD}	Energy output during the i^{th} interval of the shutdown ramp process [MWh]
E_{gsi}^{SU}	Energy output during the i^{th} interval of the startup ramp process type s [MWh]
P_{gi}^{SD}	Power output at the beginning of the i^{th} interval of the shutdown ramp process [MW]
P_{gsi}^{SU}	Power output at the beginning of the i^{th} interval of the startup ramp process type s [MW].
RD_g	Ramp-down capability [MW/h]
RU_g	Ramp-up capability [MW/h]
RU_g	Shutdown ramping capability [MW/h]
SU_g	Startup ramping capability [MW/h]
SD_g^{D}	Duration of the shutdown process [h]
SU_{gs}^{D}	Duration of the startup process type s [h]
T_{gs}^{SU}	Time defining the interval limits of the startup segment s, $[T_{gs}^{SU}, T_{g,s+1}^{SU})$ [h]
SU_g	Minimum down time [h]
TU_g	Minimum up time [h].
W_t^{E}	Available renewable energy for hour t [MWh]
W_t^{P}	Available renewable power at end of hour t [MW]

Continuous Non-Negative Variables

w_t^{E}	Renewable energy output for hour t [MWh]
w_t^{P}	Renewable power output at the end of hour t [MW]
e_{gt}	Energy output above minimum output for hour t [MWh]
$\hat{e}_{gt}$	Total energy output at the end of hour t, including startup and shutdown trajectories [MWh]
p_{gt}	Power output above minimum output at the end of hour t [MW]
$\hat{p}_{gt}$	Total power output at the end of hour t, including startup and shutdown trajectories [MW]

Binary Variables

u_{gt}	1 if the unit is producing above minimum output and 0 otherwise
y_{gt}	1 if the unit starts up and 0 otherwise
z_{gt}	1 if the unit shuts down and 0 otherwise
δ_{gst}	Startup type s. 1 if the unit starts up and has been previously down within $[T_{gs}^{SU}, T_{g,s+1}^{SU})$ hours

A.1. Energy-based UC formulation

The energy-based formulation, based on Gentile et al. [20], is the traditional approach to unit commitment. The objective is to minimise the total cost of production:

$$\min \sum_{t \in \mathcal{T}} \left(\sum_{g \in \mathcal{G}} \left[C_g^{\mathrm{NL}} u_{gt} + \sum_{s \in S_g} C_{gs}^{\mathrm{SU}} \delta_{gst} + C_g^{\mathrm{SD}} z_{gt} \right] + \sum_{g \in \mathcal{G}} C_g^{\mathrm{LV}} \hat{e}_{gt} + C^{\mathrm{VW}} w_t^{\mathrm{E}} \right) \tag{A.1}$$

subject to an energy balance

$$\sum_{g \in \mathcal{G}} \hat{e}_{gt} = D_t^{\mathrm{E}} - w_t^{\mathrm{E}} \qquad \forall\, t \tag{A.2}$$

Individual generators must adhere to commitment and startup/shutdown logic:

$$u_{gt} - u_{g,t-1} = y_{gt} - z_{gt} \qquad \forall\, g, t \tag{A.3}$$

$$\sum_{i=t-\mathrm{TU}_g+1}^{t} y_{gi} \leqslant u_{gt} \qquad \forall\, g, t \tag{A.4}$$

$$\sum_{j=t-\mathrm{TU}_g+1}^{t} z_{gi} \leqslant 1 - u_{gt} \qquad \forall\, g, t \tag{A.5}$$

The startup trajectories differ per startup type, which depends on the amount of time the generator has been offline:

$$\delta_{gst} \leqslant \sum_{i=T_{gs}^{\mathrm{SU}}}^{T_{g,s+1}^{\mathrm{SU}}-1} z_{g,t-i} \qquad \forall\, g, s \in [1, S_g), t \tag{A.6}$$

$$\sum_{s \in S_g} \delta_{gst} = y_{gt} \qquad \forall\, g, t \tag{A.7}$$

Generators are also subject to maximum electricity production limits, and limits on ramping during the startup and shutdown phases.

$$e_{gt} \leqslant (\overline{P}_g - \underline{P}_g) u_{gt} - (\overline{P}_g - SD_g) z_{g,t+1} - \max(SD_g - SU_g, 0) y_{g,t} \qquad \forall\, g \in \mathcal{G}^1, t \tag{A.8}$$

$$e_{gt} \leqslant (\overline{P}_g - \underline{P}_g) u_{gt} - (\overline{P}_g - SU_g) y_{gt} - \max(SU_g - SD_g, 0) z_{g,t+1} \qquad \forall\, g \in \mathcal{G}^1, t \tag{A.9}$$

$$e_{gt} \leqslant (\overline{P}_g - \underline{P}_g) u_{gt} - (\overline{P}_g - SU_g) y_{gt} - (\overline{P}_g - SD_g) z_{g,t+1} \qquad \forall\, g \notin \mathcal{G}^1, t \tag{A.10}$$

The change in electricity output is constrained by ramping limits:

$$-RD_g \leqslant e_{gt} - e_{g,t-1} \leqslant RU_g \qquad \forall\, g, t \tag{A.11}$$

The total amount of electricity production for thermal generators is the sum of the startup trajectory, the shutdown trajectory, and the output when up; renewable generators generate between 0 and the maximum available.

$$\hat{e}_{gt} = \sum_{s=1}^{S_g} \sum_{i=1}^{SU_{gs}^{\mathrm{D}}} E_{gsi}^{\mathrm{SU}} \delta_{gs,(t-i+SU_{gs}^{\mathrm{D}}+2)} + \sum_{i=2}^{SD_g^{\mathrm{D}}+1} E_{gi}^{\mathrm{SD}} z_{g,(t-i+2)} + \tag{A.12}$$

$$+ \underline{P}_g (u_{gt} + y_{g,t+1}) + e_{gt} \qquad \forall\, g \in \mathcal{G}^{\mathrm{S}}, t$$

$$0 \leqslant w_t^{\mathrm{E}} \leqslant W_t^{\mathrm{E}} \qquad \forall\, t \tag{A.13}$$

The startup and shutdown trajectories are linear interpolations of reaching minimal stable output from 0, and vice versa. Note that these are derived from the power output at the beginning of the respective segment, and the energy output

$$E_{gsi}^{\mathrm{SU}} = \frac{P_{gsi}^{\mathrm{SU}} + P_{gs(i+1)}^{\mathrm{SU}}}{2} \qquad \forall\, g, s \tag{A.14}$$

$$E_{gi}^{\mathrm{SD}} = \frac{P_{gi}^{\mathrm{SD}} + P_{g(i+1)}^{\mathrm{SD}}}{2} \qquad \forall\, g \tag{A.15}$$

where

$$P_{gsi}^{\mathrm{SU}} = \frac{\underline{P}_g}{SU_{gs}^{\mathrm{D}}} \cdot (i-1) \qquad \forall\, g, s \tag{A.16}$$

$$P_{gi}^{\mathrm{SD}} = \underline{P}_g - \frac{\underline{P}_g}{SD_g^{\mathrm{D}}} \cdot (i-1) \qquad \forall\, g \tag{A.17}$$

A.2. Power-based UC formulation

The power-based UC formulation is based on [21].
The objective is to minimise the total cost of production:

$$\min \sum_{t \in \mathcal{T}} \left(\sum_{g \in \mathcal{G}} \left[C_g^{\mathrm{NL}} u_{gt} + \sum_{s \in \mathcal{S}_g} C_{gs}^{\mathrm{SU}'} \delta_{gst} + C_g^{\mathrm{SD}'} z_{gt} \right] + \sum_{g \in \mathcal{G}} C_g^{\mathrm{LV}} \hat{e}_{gt} + C^{\mathrm{VW}} w_t^{\mathrm{E}} \right)$$

(A.18)

where once again

$$C_{gs}^{\mathrm{SU}'} = C_{gs}^{\mathrm{SU}} + C_g^{\mathrm{NL}} SU_{gs}^{\mathrm{D}} \qquad \forall\, g, s$$

(A.19)

$$C_g^{\mathrm{SD}'} = C_g^{\mathrm{SD}} + C_g^{\mathrm{NL}} SD_g^{\mathrm{D}} \qquad \forall\, g$$

(A.20)

The key difference between the two formulations is that the supply–demand balance is written in terms of power.

$$\sum_{g \in \mathcal{G}} \hat{p}_{gt} = D_t^{\mathrm{P}} - w_t^{\mathrm{P}} \qquad \forall\, t$$

(A.21)

As above, individual units must adhere to commitment constraints, as well as their start-up and shut-down logic:

$$u_{gt} - u_{g,t-1} = y_{gt} - z_{gt} \qquad \forall\, g, t$$

(A.22)

$$\sum_{i=t-\mathrm{TU}_g+1}^{t} y_{gi} \leqslant u_{gt} \qquad \forall\, g, t$$

(A.23)

$$\sum_{j=t-\mathrm{TU}_g+1}^{t} z_{gi} \leqslant 1 - u_{gt} \qquad \forall\, g, t$$

(A.24)

The power output for slow-start units is again the sum of output in a shutdown or startup trajectory and the output in the 'on'-state:

$$\hat{p}_{gt} = \sum_{s=1}^{S_g} \sum_{i=1}^{SU_{gs}^{\mathrm{D}}} P_{gsi}^{\mathrm{SU}} \delta_{gs,(t-i+SU_{gs}^{\mathrm{D}}+2)} + \sum_{i=2}^{SD_g^{\mathrm{D}}+1} P_{gi}^{\mathrm{SD}} z_{g,(t-i+2)} +$$

$$+ \underline{P}_g(u_{gt} + y_{g,t+1}) + p_{gt} \qquad \forall\, g \in \mathcal{G}^{\mathrm{S}}, t$$

(A.25)

There are different startup trajectories for different startup types, depending on for how long the generator has been offline.

$$\delta_{gst} \leqslant \sum_{i=T_{gs}^{\mathrm{SU}}}^{T_{g,s+1}^{\mathrm{SU}}-1} z_{g,t-i} \qquad \forall\, g, s \in [1, S_g), t$$

(A.26)

$$\sum_{s \in \mathcal{S}_g} \delta_{gst} = y_{gt} \qquad \forall\, g, t$$

(A.27)

The startup and shutdown trajectories are again linear interpolations.

$$P_{gsi}^{\mathrm{SU}} = \frac{\underline{P}_g}{SU_{gs}^{\mathrm{D}}} \cdot (i-1) \qquad \forall\, g, s$$

(A.28)

$$P_{gi}^{\mathrm{SD}} = \underline{P}_g - \frac{\underline{P}_g}{SD_g^{\mathrm{D}}} \cdot (i-1) \qquad \forall\, g$$

(A.29)

Since the total output of generators is defined in terms of power, rather than energy, we must write the ramping constraints and resulting energy output as a function of the power output:

$$-RD_g \leqslant p_{gt} - p_{g,t-1} \leqslant RU_g \qquad \forall\, g, t$$

(A.30)

$$\hat{e}_{gt} = \frac{\hat{p}_{g,t-1} + \hat{p}_{gt}}{2} \qquad \forall\, g, t$$

(A.31)

Similarly, the wind output in power and energy, respectively, are

$$w_t^{\mathrm{P}} \leqslant W_t^{\mathrm{P}} \qquad \forall\, t$$

(A.32)

$$w_t^{\mathrm{E}} = \frac{w_{t-1}^{\mathrm{P}} + w_t^{\mathrm{P}}}{2} \qquad \forall\, t$$

(A.33)

Finally, we also express the startup and shutdown output in power.

$$p_{gt} \leqslant (\overline{P}_g - \underline{P}_g) u_{gt} - (\overline{P}_g - SD_g) z_{g,t+1} + (SU_g - \underline{P}_g) y_{g,t+1} \qquad \forall\, g, t$$

(A.34)

References

[1] Morales-España Germán, Ramírez-Elizondo Laura, Hobbs Benjamin F. Hidden power system inflexibilities imposed by traditional unit commitment formulations. Appl Energy 2017;191:223–38.

[2] Philipsen Rens, Morales-España Germán, de Weerdt Mathijs, de Vries Laurens. Imperfect Unit Commitment decisions with perfect information: A real-time comparison of energy versus power. In: 19th Power systems computation conference, PSCC 2016; 2016.

[3] NERC. 22:00 Frequency Excursions (Final Report). Technical report; 2002.

[4] ENTSO-E and Eurelectric. Deterministic frequency deviations - root causes and

proposals for potential solutions. Technical report December; 2011.

[5] UCTE. Frequency quality investigation, excerpt of the final report. Technical report; 2008.

[6] Abdul-Rahman Khaled H, Alarian Hani, Rothleder Mark, Ristanovic Petar, Vesovic Bogdan, Lu Bo. Enhanced system reliability using flexible ramp constraint in CAISO market. In: IEEE power and energy society general meeting; 2012.

[7] Navid N, Rosenwald G. Market solutions for managing ramp flexibility with high penetration of renewable resource. IEEE Trans Sustain Energy 2012;3(4):784–90.

[8] Neuhoff Karsten, Ritter Nolan, Salah-Abou-El-Enien Aymen, Vassilopoulos Philippe. Intraday markets for power: discretizing the continuous trading. DIW Berlin, (March 2016); 2016.

[9] Deane JP, Drayton G, Ó Gallachóir BP. The impact of sub-hourly modelling in power systems with significant levels of renewable generation. Appl Energy 2014;113:152–8.

[10] Pandzic Hrvoje, Dvorkin Yury, Wang Yishen, Qiu Ting, Kirschen Daniel S. Effect of time resolution on unit commitment decisions in systems with high wind penetration. In: 2014 IEEE PES general meeting—conference & exposition; 2014.

[11] Wang Jiadong, Wang Jianhui, Liu Cong, Ruiz Juan P. Stochastic unit commitment with sub-hourly dispatch constraints. Appl Energy 2013;105:418–22.

[12] Bakirtzis Emmanouil A, Biskas Pandelis N, Labridis Dimitris P, Bakirtzis Anastasios G. Multiple time resolution unit commitment for short-term operations scheduling under high renewable penetration. IEEE Trans Power Syst 2014;29(1):149–59.

[13] vom Stein Denis, van Bracht Niklas, Maaz Andreas, Moser Albert. Development of adaptive time patterns for multi-dimensional power system simulations. In: 2017 14th International Conference on the European Energy Market (EEM); 2017. p. 1–5

[14] Matus Marcelo, Cceres Nicols, Püschel-Lvengreen Sebastin, Moreno Rodrigo. Chebyshev based continuous time power system operation approach. 2015;1:1–5.

[15] Morales-España Germán, Ramos Andrés, García-González Javier. An MIP formulation for joint market-clearing of energy and reserves based on ramp scheduling. IEEE Trans Power Syst 2014;29(1):476–88.

[16] Yang Yuanchao, Wang Jianhui, Guan Xiaohong, Zhai Qiaozhu. Subhourly unit commitment with feasible energy delivery constraints. Appl Energy 2012;96:245–52.

[17] Neon. DIW, Europa-Universität Flensburg, and Technical University of Berlin. Open Power System Data; 2016.

[18] Hidalgo Gonzalez Ignacio, Sylvain Quoilin, Zucker Andreas. Dispa-SET 2.0: unit commitment and power dispatch model. Technical report, Joint Research Center, Petten; 2014, 10.2790/399921.

[19] Schröder Andreas, Kunz Friedrich, Meiss Jan, Mendelevitch Roman, von Hirschhausen Christian. Current and prospective costs of electricity generation until 2050. Technical Report 68, Deutsches Institut fuer Wirtschaftsforschung, Berlin; 2013.

[20] Gentile C, Morales-España G, Ramos A. A tight MIP formulation of the unit commitment problem with start-up and shut-down constraints. EURO J Comput Optimiz 2017;5(1-2):177–201.

[21] Morales-España Germán, Gentile Claudio, Ramos Andrés. Tight MIP formulations of the power-based unit commitment problem. OR Spectrum 2015;37(4):929–50.

[22] Elia. Implementatie van de day-ahead marktkoppeling in de regio NWE (Noord-West Europa) Overzichtsnota. Technical report; 2013.

[23] EPEX SPOT. EUPHEMIA Public Description. Technical report; 2015.

[24] O'Neill Richard P, Sotkiewicz Paul M, Hobbs Benjamin F, Rothkopf Michael H, Stewart William R. Efficient market-clearing prices in markets with nonconvexities. Eur J Oper Res 2005;164(1):269–85.

[25] Myerson Roger B, Satterthwaite Mark A. Efficient mechanisms for bilateral trading. J Econ Theory 1983;29:265–81.

[26] Groves Theodore. Incentives in teams. Econometrica: J Econ Soc 1973;41(4):617–31.

[27] Green Jerry, Laffont Jean-Jacques. Characterization of satisfactory mechanisms for the revelation of preferences for public goods. Econometrica 1977;45(2):427–38.

[28] 50 Hertz. ENTSO-E Netzfrequenz; 2017.