Imperfect Unit Commitment Decisions with Perfect Information: a Real-time Comparison of Energy versus Power

Philipsen, Rens; Morales Espana, German; de Weerd, Mathijs; de Vries, Laurens

DOI

[10.1109/PSCC.2016.7540923](https://doi.org/10.1109/PSCC.2016.7540923)

Publication date

2016

Document Version

Accepted author manuscript

Published in

19th Power Systems Computation Conference, PSCC 2016

Citation (APA)

Philipsen, R., Morales Espana, G., de Weerd, M., & de Vries, L. (2016). Imperfect Unit Commitment Decisions with Perfect Information: a Real-time Comparison of Energy versus Power. In *19th Power Systems Computation Conference, PSCC 2016* (pp. 1-7). IEEE .
<https://doi.org/10.1109/PSCC.2016.7540923>

Important note

To cite this publication, please use the final published version (if applicable).
Please check the document version above.

Copyright

Other than for strictly personal use, it is not permitted to download, forward or distribute the text or part of it, without the consent of the author(s) and/or copyright holder(s), unless the work is under an open content license such as Creative Commons.

Takedown policy

Please contact us and provide details if you believe this document breaches copyrights.
We will remove access to the work immediately and investigate your claim.

Imperfect Unit Commitment Decisions with Perfect Information: a Real-time Comparison of Energy versus Power

Rens Philipsen, Germán Morales-España, Mathijs de Weerd, Laurens de Vries

Abstract—In order to cope with fluctuations and uncertainty, power systems rely on contracted reserves. The day-ahead Unit Commitment (UC) is the short-term planning process which is commonly used to schedule these resources at minimum cost, while operating the system and units within secure technical limits. This paper shows through the evaluation of deterministic cases that conventional energy-based UC formulations lead to inefficient use of reserves in real-time operation to deal with completely known deterministic events. These inefficient decisions are inherent to the assumptions underlying the energy-based formulation. A power-based formulation of the unit commitment problem is shown to avoid infeasible schedules at minimal cost, increasing economic efficiency and system security.

I. INTRODUCTION

Power systems worldwide are facing a significant growth of renewable electricity generation, such as from wind and sun, driven by concerns for the environment and energy security. Due to their intermittency, however, their realised hourly production does not necessarily match the day-ahead prediction. To maintain the supply-demand balance, as indicated by a constant frequency, adequate system resources (reserves) must be scheduled in advance to guarantee that the power system has the flexibility to compensate for possible variations in load and (renewable) output. As the generation mix in most power systems has been undergoing a transition in recent years, ensuring the continued availability of sufficient reserves is of prime importance [3, 12].

For sufficient reserves to be present, they must be scheduled in advance. The day-ahead scheduling process, which ensures availability of generation and reserve capacity, can be represented by a Unit Commitment (UC) formulation. UC is widely held to be theoretically optimal. Recent improvements to UC formulations have seen better inclusion of security constraints [14], ramping rates [1], start-up trajectories [9], and better computational performance [5]. Despite these recent advancements, and contrary to popular belief, traditional UC is not optimal, even with perfect information: reserves are currently also required to compensate frequency deviations occurring at the change of the hour during ramping periods. These frequency excursions exist across all power systems, and are deterministic and predictable: system operators in both the U.S. and Europe have identified market dynamics as their primary cause, rather than uncertainty, to which most frequency fluctuations are attributed [4, 11]. These events significantly impact system security and reliability: UCTE [13]

states that a generation outage of 1300MW causes a frequency dip of 50 mHz on the European synchronous zone, while market-related frequency dips in the evening exceed 100 mHz on average. This raises the challenge of designing markets which send efficient incentives that continuously align supply and demand, allowing optimal usage of the flexibility provided by both sides of the market, and more efficient usage of the entire power system in general.

Preventing these deterministic imbalances from occurring is one of the major challenges faced in power systems operation. To do so, we must have a thorough understanding of their causes: how do market dynamics lead to imbalance? Moreover, we aim to quantify the problem: what is the social cost of imperfect scheduling? The primary goal of this paper is to provide answers to these two questions. Our approach is as follows. First, we provide a theoretical comparison of energy-based and power-based UC, where we identify the root causes of the inefficiencies of energy-based scheduling. Rather than settling for a rough identification of the shortcomings of energy-based scheduling, we provide a detailed discussion of direct causes and consequences using very small, stylised examples. Second, we quantify the error of energy-based representations of demand, and show for the IEEE 118-bus test system how the different formulations perform in terms of efficiency and reserves dispatched. We show that even under the most ideal conditions, energy-based UC is unable to accurately account for the physical constraints of available generators. We stress that these inefficiencies exist even in a deterministic case; they can therefore not be attributed to the presence of uncertainty in supply or demand.

II. ENERGY-BASED VERSUS POWER-BASED UNIT COMMITMENT

Energy-based UC makes a coarse approximation of generator capabilities by modelling the output of a generator as energy levels e within a large scheduling interval (usually 1 hour) t , as displayed in (1), in MWh. Maximum up (RU^{max}) and down (RD^{max}) rates are then imposed on the difference between energy levels in subsequent hours. Additionally, energy-based UC models assume that, if a unit is turned on ($u_t = 1$), it starts and ends electricity production at its minimum output e^{min} , as in (2), ignoring the start-up and shut-down power trajectories. As a consequence, there may be a large amount of energy that is not allocated by the UC

Figure 1: Network topology, with generator G1 separated from the rest of the network by line L1. Generators are described by $[P^{max}, RU^{max}]$

but which is present during real time operation, thus affecting the total load balance.

$$-RD^{max} \leq e_t - e_{t-1} \leq RU^{max} \quad (1)$$

$$u_t \cdot e^{min} \leq e_t \leq e^{max} \cdot u_t \quad (2)$$

These constraints are typically used to match the hourly energy-based market. Physically, however, ramp constraints are imposed on instantaneous power output rather than average output.

To better align markets and operation, a number of possible approaches exist. Weissbach and Welfonder [15] suggest a uniform ramp rate for generators to reduce frequency deviations. A different approach would be to reformulate the scheduling problem to better account for physical constraints, rather than to focus on energy output only. The power-based UC uses a piecewise-linear representation of thermal generation, load, and wind, instead of the conventional stepwise energy profiles [7]. These power trajectories represent the instantaneous electricity consumption (or production), rather than averaging across the entire hour into a single energy block. This requires the use of a different time base: variable t no longer refers to an interval, but to a single point in time (e.g., the end of such an interval). Observe the power-based equivalents of (1) and (2). Equation (3) sets ramp constraints on the instantaneous power output P in MW, representing the maximum difference in power generation between two moments in time, instead of a difference in total energy output between two time intervals. Equations (4) and (5) constrain the total power output $\hat{P}_{g,t}$, whenever the generator is operating above minimum output, as a function of its commitment, start-up, and shut-down state (respectively u, v, w ; see [7] for details).

$$-RD^{max} \leq P_t - P_{t-1} \leq RU^{max} \quad (3)$$

$$P_t \leq (P^{max} - P^{min}) \cdot (u_t - w_{t+1}) \quad (4)$$

$$\hat{P}_t = P^{min}(u_t + v_{t+1}) + P_t \quad (5)$$

III. THEORETICAL ANALYSIS

In order to clearly identify which differences in real-time operation arise from this reformulation into power trajectories,

Figure 2: Possible power trajectories matching the energy schedule from $t = 0$ to $t = 3$.

we analyse a number of situations using a stylised example. This example system is used to execute day-ahead schedules as determined by different formulations of the unit commitment problem. Three cases are considered; the myopic execution of an energy-based schedule (pure self-dispatch); the real-time dispatch of an energy-based schedule (5-minute interval, as in ERCOT); and myopic execution of a power-based schedule. Our example network is set up as shown in Figure 1, which also shows the data associated with different instances of the units and the transmission line, upon which the discussion in this section is based. Our analysis is based on a program time unit (PTU) or time step size of 1 hour, for an easier calculation of the produced electricity. The results, of course, also hold for shorter PTUs, although it must be noted that the shorter the interval, the more precise a stepwise representation will be.

A. Multiple solutions to day-ahead energy schedule

Matching hourly energy schedules can result in a mismatch on a second-to-second basis. Figure 2 shows the energy bids submitted by the consumer. An energy-based market would reveal nothing more than that demand in period 0-1 is 50 MWh, 100 MWh in 1-2, and 150MWh from $t = 2$ onwards. This is already a very coarse approximation: in reality, demand follows a smooth curve, without the discontinuous jumps at the change of the hour. The power trajectories are piecewise linear, still approximating actual demand but doing so less coarsely.

The power trajectories, which a generator can follow to supply this energy block, follow directly from its output at $t = 0$. Figure 2 also shows a number of possible power trajectories for generator G_1 , not all of which necessarily feasible; trajectory C, for example, requires a 200 MW ramp-up in $t = 1 \rightarrow 2$, even though the energy output increases by only 50 MWh compared to the previous PTU. Many more power schedules (including those with changing ramps during the PTUs), all matching the energy bids, can of course be imagined.

Meanwhile, the power trajectory of demand is completely unknown to us, as the market result only tells us what the energy consumption is. This results in a mismatch between

Figure 3: Power trajectories and their associated energy output for G_2 , given an output of 50 MWh in PTU $0 \rightarrow 1$

supply and demand on a second-to-second time scale, which is rather problematic: one of the core purposes of the electricity market, after all, is to align supply and demand as to ensure a continuous balance between the two.

The possibility of following different power schedules to meet an energy schedule has a number of implications for how reserves are deployed, balancing payments are made, and how congestion is managed. We examine these issues using the example introduced above, making some slight modifications where necessary.

B. Imperfect inclusion of ramping constraints

Since there are multiple power trajectories possible to implement an energy schedule, correctly including the ramping constraints in an energy-based day-ahead planning is a difficult job. Ramping constraints are, after all, physically expressed as the change in power output between two individual points in time, rather than the change in energy output between two intervals. To illustrate, consider generator G_2 , with $RU_{\max} = 50$ (intentionally omitting its unit). Figure 3 displays the differences between formulating this as an energy constraint, vis-a-vis formulating it as a power constraint. All power trajectories (lines) result in an energy output of 50 MWh in the first PTU, as shown by the associated energy levels (bars). Formulating the ramping constraint in energy, the upper bound on the energy output in the next PTU is 100 MWh.

As we have shown above, however, a number of power trajectories can fulfil the energy demand. We display three power profiles, all fulfilling the energy demand in the first PTU. Of the three profiles shown, only the solid line can provide the ramping rate assumed by the energy-based schedule. The other bars also show the energy output for their associated power trajectories. In the worst case, the energy output in the second PTU cannot change at all, compared to the first.

The power output at the start of the PTU accurately represents the state of the generator. Expressing the ramping capability in terms of power, $P_{t+1} \leq P_t + RU_{\max}$ (MW). Energy output during the previous PTU, on the other hand, is not a good indicator of a generator's possible production during the next PTU: it overestimates flexibility. Ramping up, the maximum energy output during e_{t+1} could be as low as

Figure 4: Day-ahead schedule based on energy bids. The power trajectories for G_1 and G_2 are cumulative.

$e_t + 0.5 \cdot (RU_{\max} - RD_{\max})$ (MWh) (if the generator was ramping down at maximum speed), or as high as $e_t + RU_{\max}$. Energy-based planning consistently assumes the latter is the case. Consequently, the error in energy-based scheduling may be as large as $0.5 \cdot RD_{\max} \cdot 1[h]$ (MWh). The same error, of course, exists for ramping down, but with RU/RD reversed. The amount of balancing energy required would then be equal to $(P_1 - P_0) \cdot 1[h]$ (MWh).

C. Inefficient use of reserves

An energy schedule which is optimal in day-ahead planning may not be optimal when it is executed in real time. Consider the case of generators G_1, G_2, G_3 from Figure 1, all producing 0 MW before $t = 0$. Matching the energy blocks in the day-ahead market, G_1 is scheduled to provide (50,100,100) MW, G_2 provides (0,0,50) MW, and G_3 remains unused.

We first examine what happens if we try to execute this schedule, as displayed in Figure 4. G_1 ramps up at 100 MW/h, delivering the contracted 50 MWh, but requiring a total of 25 MWh of both up- and down-reserves to maintain the supply-demand balance; until halfway through it is not delivering enough, but in the second half of the PTU there is overproduction. Continuing to PTU 1-2, the generator can deliver 100 MWh by producing continuously at full load, while demand ramps up from 50 MW to 150 MW; once again, we require reserves throughout the entire period, as in the first half there is overproduction, and in the second there is underproduction. During the third period, G_2 is scheduled to supply 50 MWh, while it can ramp at 50 MW/h. Starting its output at 0 MW, G_2 can only supply 25 MWh, the case as described in the previous section.

We now examine the case in which we use real-time (every 5 minutes) dispatching to follow the demand trajectory, as shown on the left in Figure 5. We dispatch G_1 and G_2 at $t = 0$, to ramp up as fast as possible. They reach a combined output of 50 MW at $t = 0.33$, while the expensive but infinitely flexible G_3 ramps up to 50 MW immediately, reducing its output to 0 as the other two ramp up. As G_1 is cheaper, G_2 also steadily reduces its output, reaching 0 MW by $t = 0.5$. At $t = 1$, G_1 is dispatched to ramp up even further, reaching its maximum output of 100 MW at $t = 1.5$. At this point, G_2 must

Figure 5: Comparing energy-based scheduling with real-time dispatch to power-based scheduling.

increase its output to keep up with the still-rising demand, but once again it can only ramp at 50 MW/h. Consequently, it can only reach 25 MW by $t = 2$, and therefore G_3 must also be dispatched at $t = 1.5$ to ramp at 50 MW/h in order to follow the demand trajectory. From $t = 2$ onwards, demand remains constant, and G_2 continues to ramp up while G_3 ramps down at the same speed. By $t = 2.5$, G_1 produces 100 MW, G_2 50 MW, and G_3 0 MW.

On the right in Figure 5, we show the results for power-based unit commitment. We still need G_3 at $t = 0$ to supply the required ramp, and use the same schedule until $t = 1$, at which point we dispatch both G_1 and G_2 to ramp up at 50 MW/h. G_1 will now reach its maximal output by $t = 2$, and can therefore supply the required ramp rate. Consequently, G_3 is no longer required to support G_2 in supplying the ramping rate after G_1 has reached its maximum output. Expressed in terms of energy, G_2 produces an additional 6.25 MWh of electricity between $t = 1$ and $t = 1.5$, 12.5 MWh between $t = 1.5$ and $t = 2$, and 6.25 MWh between $t = 2$ and $t = 2.5$. While we require an additional 25 MWh from G_2 , we require 12.5 MWh fewer from G_3 between $t = 1.5$ and $t = 2.5$. Since a quick-start unit such as G_3 tends to be far more expensive than a slower-ramping unit, scheduling based on power trajectories reduces the total cost of operation, and frees up reserves for other purposes, reducing the cost of scheduling reserves as well.

An upper bound on the amount of balancing energy required, assuming constant ramping speed of both generator and load (with matching energy schedules!), would be reached if they were to show exact opposite behaviour (e.g. trajectories 1 and 3 in Figure 3). The amount of balancing energy required would then be equal to $(P_1 - P_0) \cdot 1[h]$ (MWh).

D. Violating transmission constraints

To investigate the effect of energy-based market clearing on the use of line capacity, we replace line L_1 by L'_1 , G_1 by G'_1 , and set its initial output to 50 MW. The energy bid by G'_1 cannot be fulfilled due to the capacity constraint, but this does not show in the day-ahead energy schedule: as before, G'_1 is set to deliver (50,100,100), while G_2 is set to deliver (0,0,50). Given its initial output, G'_1 's power profile must be 50 $\rightarrow$ 50, 50 $\rightarrow$ 150, 150 $\rightarrow$ 50, violating the capacity constraint on the line. Real-time dispatch would disallow the increase beyond 100 MW halfway through the

second PTU. Essentially, we see the same problem as in the previous section; since the congested line prevents G_1 from meeting the total demand, a downstream generator (G_2 or G_3) must step in. Under real-time dispatch, the downstream generators would be dispatched halfway through $1 \rightarrow 2$. The cheaper of the two, G_2 , cannot supply the required ramp, and therefore the expensive G_3 is also needed. A more efficient solution would be to take the limited ramping rate of G_2 into consideration at an earlier stage and start ramping at $t = 1$, reducing the ramping speed of G_1 to have it reach 100 MW by the end of $1 \rightarrow 2$, rather than halfway through. This results in the generators precisely matching the demand trajectory, without requiring the expensive G_3 to step in.

We foresee similar problems on an interconnector between two price zones. Assume the price difference changes sign at $t = 1$. In PTU $0 \rightarrow 1$, users want to maximise the energy exchange in one direction (maximising profit), but in $1 \rightarrow 2$ want to completely reverse the flow. As before, the profit-maximising power profile matching the energy bids is not constant transfer at the maximum average rate, but sees an increase above the line capacity starting at $t = 0$, then a sudden ramp down to drop below line capacity at $t = 1$, ensuring the total energy transfer does not exceed the hourly energy capacity. From $t = 1$ the ramp down is continued until the line capacity in the other direction is exceeded, followed by a ramp up to reach the capacity limit in power (at which it can then steadily transfer during the next hours). In total the power capacity of the interconnector is exceeded once in both directions, as parties maximise the energy exchange in order to fulfil the accompanying trades. This capacity violation, assuming equal up- and down rates, can reach up to half the line capacity.

E. Inefficient commitment decisions

In the previous examples, we assumed that the necessary reserves are always available. Spinning reserves, however, require a generator to be up and running, and therefore its use must be accounted for in the day-ahead schedule. Consider the case described in Figure 5, but assume for the moment that G_2 is only able to supply its indicated ramp rate if it producing above some minimum output, and ramps more slowly if it is in its start-up phase. As G_2 is scheduled to start producing at $t = 2$, its start-up planning will be made with this in mind, and

Figure 6: Demand curve, and representations by interpolated piecewise linear curve and stepwise blocks

the unit will not be available to provide its full ramp starting at $t = 0$, nor at $t = 1$. In the reverse case, more spinning reserve capacity may be scheduled than would be necessary under a power-based schedule, leading to more thermal generators running below full capacity, thus reducing the efficiency.

IV. NUMERICAL EVALUATION

In this section we numerically compare the performance of the power-based and energy-based unit commitment formulations in a deterministic scenario. First, we compare how well different approximations can represent actual demand. Figure 6 shows an actual demand curve for a normal day, along with two approximations. The stepwise energy blocks are unable to correctly account for the ramps at the beginning and end of the day, with a total error of 1429.95 MWh, 1.3% of total demand. An interpolated linear function, while not a perfect match, is able to follow demand much more closely, with an error of 346.27 MWh, or 0.32%. This once again shows energy blocks are a very coarse method for approximating instantaneous energy demand.

But how does that translate to errors in day-ahead scheduling? For this purpose, we use the modified IEEE 118-bus test system described in [8] for a time span of 24 hours. The system has 118 buses, 186 transmission lines, 91 loads, 64 thermal units, 10 of which are quick-start units (i.e., units that can produce from 0 to above the minimum output in 5 minutes), and three buses with wind production. All system conditions are deterministic; that is, load and wind are assumed to be perfectly known.

Three different UC formulations are implemented:

- En-UC: traditional energy-based UC, in which SU and SD trajectories are not included in the day-ahead scheduling stage, although they actually do occur in real time.
- EnSUSD-UC: traditional energy-based UC, including startup and shutdown power trajectories in both day-ahead scheduling and real-time operations.
- Pw-UC: the power-based UC, which by definition includes startup and shutdown power trajectories from the day-ahead.

As described in Section I, the deterministic En-UC is the commonly used UC approach, in which the energy demand is represented using energy levels (hourly-averaged generation)

in a stepwise fashion over time. All constraints involving generation levels are applied to these energy levels. For this study, we use the UC formulation in [6] to represent En-UC. The startup and shutdown trajectories are included in EnSUSD-UC using the model in [9].

The deterministic power-based UC proposed in [7, 10] draws a clear distinction between power and energy. Demand and generation are modelled as hourly piecewise-linear functions representing their instantaneous power trajectories. The schedule of generating unit output is no longer an energy stepwise function, but a smoother piece-wise power function.

Since all UC formulations are evaluated under perfect information, no reserve requirements are imposed. All formulations, however, include demand balance and transmission constraints, and the following units constraints: minimum up and down times, generating limits, ramp limits, and, startup and shutdown power trajectories, where different startup costs and startup power trajectories are modeled for EnSUSD and Pw depending on how long the units have been off-line. All the models were carried out using CPLEX 12.6.1 on an Intel-Xeon (64-bit) 3.7-GHz personal computer with 16 GB of RAM memory. The problems are solved until they hit a time limit of 2 hours or until they reach an optimality tolerance of 0.05% (none of the UC problems exceeded the time limit).

To assess the performance of the different scheduling approaches, we make a clear difference between the scheduling stage and the real-time dispatch stage. In the *scheduling stage*, the different UC problems are solved to obtain the hourly commitment schedule for the 54 slow-start units for 24 hours. In the *real-time dispatch stage*, the (hourly) commitment decisions resulting from the scheduling stage are fixed, and the real-time (5-min) dispatch decisions are optimized to supply demand at minimum cost for a time span of 24 hours using a network-constrained economic dispatch. The real-time dispatch model also take decisions on the 10 quick-start units, thus the dispatch stage mimics the actual real-time system operation in which generating units are dispatched to supply the demand every 5 minutes, while commitment decisions are allowed to be taken for the quick-start units every 15 minutes. This is an approximation of the California ISO market design, in which the fifteen minute market includes both dispatch and quick-start commitment decisions, and a 5 minute market is dispatch-only.

V. RESULTS

Table I shows the performance of the different deterministic UC formulations in eight different aspects, three related to the day-ahead scheduling stage, three to the real-time dispatch stage, and two comparing the day-ahead schedules with the actual dispatch. Scheduling stage: 1) Total production costs (TC) obtained from the optimal UC solution; 2) number of startups (SU); and 3) percentage of wind curtailment (Curt). Real-time dispatch stage: 4) Total production costs (TC) obtained from the optimal network-constrained dispatch/quick-start commitment solution; 5) number of startups of quick-start units (QSU); and 6) percentage of wind production curtailed

TABLE I: Performance of Different Deterministic UC policies

Deterministic UC	Scheduling (hourly)			Real-time Dispatch (5-min)			Sch vs. Rtd*	
	TC [k\$]	SU [#]	Curt [%]	TC [k\$]	QSU [#]	Curt [%]	TC rtd/sch	Curt rtd/sch
En	721.67	6	6.28	721.21	1	10.39	0.999	1.65
EnSUSD	701.70	9	6.45	712.25	1	8.44	1.015	1.31
Pw	706.44	11	7.50	710.72	0	7.96	1.006	1.06

* 'sch' denotes the scheduling stage, 'rtd' denotes the real-time dispatch stage

(Curt). The final two metrics compare the outcomes predicted by the scheduling model vs. what was actually realized in the real-time dispatch stage (Sch vs Rtd): 7) the ratio of the TC obtained from the dispatch stage to that predicted in the scheduling stage (TC rtd/sch); and 8) the ratio of the dispatch stage (actual) curtailment to that predicted by the scheduling stage (Curt rtd/sch).

From the scheduling stage in Table I, EnSUSD and Pw (UC formulations including startup and shutdown trajectories) started up more units than the traditional energy-based UC formulation En, which does not include startup and shutdown trajectories. Because En ignores the startup and shutdown trajectories, it implicitly assumes these units can start-up faster than they can in reality. EnSUSD and Pw, by contrast, introduce inflexible power trajectories during the startup and shutdown processes of the units, which need to be accommodated by other generating units thus maintaining the supply-demand balance.

The inflexibility imposed by the startup and shutdown trajectories leads EnSUSD and Pw to curtail more wind than En. Furthermore, Pw presents the highest curtailment due to the inflexibility introduced by modeling ramp constraints in power, which tend to be underestimated, unlike En and EnSUSD which tend to overestimate the units' ramping capabilities.

Notice that En presents the highest production costs (TC), this is mainly because En accounts the costs induced by the energy produced during the startup and shutdown processes while ignoring the energy itself in the schedule.

In terms of overall economic efficiency, and only based on the scheduling stage results, EnSUSD presents a better performance (TC) than Pw (highest TC). This is mainly because EnSUSD tends to overestimate the units ramping capabilities, therefore scheduling less resources than Pw, resulting in a lower cost. These ex-ante results are commonly used to define the economic efficiency of a scheduling approach (e.g., [2]). In reality, we should be concerned with the real-time operation results, as evaluation using the results from day-ahead scheduling implicitly assumes the day-ahead planning is executed. As we've shown in previous sections, this assumption does not hold even under deterministic conditions. We therefore examine the results that follow from the real-time dispatch.

From the real-time dispatch stage in Table I, which shows the actual system operation and performance, the results between EnSUSD and Pw are completely opposite to those of the scheduling stage. Now, EnSUSD shows higher operating costs (TC) and curtailment than Pw. Regarding curtailment performance, En was supposed to accommodate the highest amount of wind (lowest Curt in the scheduling stage) and

Pw the lowest; real-time system operation shows completely opposite results, where En presents the highest curtailment and Pw the lowest.

Comparing the expected scheduling values with the actual real-time dispatch values, the traditional energy-based UC scheduling approach En presented an actual curtailment of 65% higher than expected. As discussed in Section III, this big difference between the scheduling and real-time stages is a natural consequence of ignoring startup and shutdown trajectories as well as scheduling energy. The startup and shutdown trajectories of generating units are inherently present in real-time operation [9]; hence, to keep the supply-demand balance, this generation, which was ignored in the scheduling stage, must be accommodated by decreasing the generating output of other units, including wind. En also needed the support of quick-start units in one occasion to supply the demand. Therefore, wind curtailment is being used as an extra flexibility source to deal with the approximations introduced by the scheduling approach. Notice that when the startup and shutdown power trajectories are included in the energy-based scheduling approach (EnSUSD) the actual wind curtailment was 31% higher than expected. EnSUSD also needed the support of quick-start units (in 1 occasion) to supply the demand. The power-based scheduling approach Pw, which also includes startup and shutdown trajectories, presents the lowest wind curtailment deviations compared with what was expected (6% higher) and did not needed the support quick-start units.

VI. CONCLUSION

We have shown that even in fully deterministic situations, traditional UC formulations fail to accurately schedule generators. A step-by-step analysis of the connection between day-ahead planning and real-time operation reveals a number of discrepancies which are incorrectly captured by existing markets, and the corresponding UC formulations. To begin with, they misrepresent demand: the hourly day-ahead energy schedule consists of stepwise, hourly aggregated energy consumption. Even with perfect information, on a normal day, this induces an error in real-time predictions (more than 1% of total daily electricity demand), incurring significant additional costs correct this mismatch in real time.

Secondly, since there are different possible power trajectories, power systems overly rely on reserves to compensate for imbalances created by the market (possibly up to $(P_1 - P_0) \cdot 1[h]$ (MWh)).

Thirdly, traditional UC formulations consistently underestimate the physical limitations imposed by transmission constraints (violations up to $0.5 \cdot l_c$ (MW)) and ramping constraints (overestimating per-hour ramp by $0.5 \cdot RD_{\max}$ (MW)), and fail to account for the energy supplied during start-up and shut-down phases. This may damage lines or require stricter capacity limits to avoid damage, increase redispatching costs, or lead to infeasible schedules if not enough capacity is available to redispatch. Moreover, in our experiments, we have shown the power-based formulation outperforms both energy-based formulations. Based on comparing the scheduling and

real-time dispatch stages for the three different UC formulations of the nominal case, we can conclude the following:

- 1) Deterministic UCs: it is important to highlight that the real-time dispatch stage uses the same information used in the scheduling stage. Therefore, the natural assumption is that the deterministic UCs result in optimal performance during the real-time dispatch. Due to this assumption, it is not common to perform a real-time evaluation and conclusions are usually drawn just based on the results obtained from the scheduling stage. However, the deterministic UC formulations are not able to face perfectly known deterministic conditions, thus causing an inefficient use of resources to manage events that were ignored in the scheduling stage.
- 2) Startup and shutdown trajectories: by ignoring these trajectories, the traditional energy-based UC presumes of a level of flexibility that the units do not have. This leads to high levels of curtailment taken as security measurement to keep the supply-demand balance.
- 3) Energy-based scheduling: even after including the startup and shutdown trajectories in the energy-based UC, the real-time dispatch shows that there is still wind curtailment about 30% higher than in the scheduling stage. Again, the traditional energy-based scheduling approach tends to overestimate the ramping capabilities of the units, hence this approach cannot guarantee that the resulting energy schedule is feasible in real-time operation. Consequently, wind curtailment and the startup of quick-start units are required in real-time operation to compensate the ramping over-estimation made in the scheduling stage.
- 4) Power-based scheduling: this approach lead to the lowest deviations in the real-time dispatch stage ($rtd/sch=1.06$). This results as a natural consequence of intra-hour variations that could not be taken into account into the hourly UC formulation.

We identify a number of avenues for future research. While these results hold for all energy-based scheduling methods (which necessarily cannot outperform UC), markets with a stronger focus on self-scheduling may face additional problems. In systems where system operators call (tertiary) reserves based on energy requirements for that PTU, the precise trajectory of these reserves cannot be controlled, leading to even more reserves being used to compensate. Meanwhile, these market inefficiencies may give rise to strategic behaviour by generators. Secondly, the precision of energy-based formulations increases as the size of the time step decreases. Using a higher resolution in day-ahead planning could therefore also be a possible solution, but it is unclear what the associated increase in computational expense would be. Additionally, in order to generalise the numerical results, additional analysis based on different power systems is required. Finally, a power-based approach to day-ahead planning requires the design of a new market structure, providing the right price signals to both generators and load.

REFERENCES

- [1] José M. Arroyo and Antonio J. Conejo. Modeling of start-up and shut-down power trajectories of thermal units. *IEEE Transactions on Power Systems*, 19(3):1562–1568, 2004.
- [2] L. Deng, B.F. Hobbs, and P. Renson. What is the Cost of Negative Bidding by Wind? A Unit Commitment Analysis of Cost and Emissions. *IEEE Transactions on Power Systems*, PP(99):1–10, 2014.
- [3] ENTSO-E. ENTSO-E Network Code on Electricity Balancing. Technical report, 2014.
- [4] ENTSO-E and Eurelectric. Deterministic frequency deviations: root causes and proposals for potential solutions. Technical Report December, 2011.
- [5] Tao Li and Mohammad Shahidehpour. Price-Based Unit Commitment : A Case of Lagrangian Relaxation Versus Mixed Integer Programming. *IEEE Transactions on Power Systems*, 20(4):2015–2025, 2005.
- [6] G. Morales-Espana, Jesus M. Latorre, and A. Ramos. Tight and Compact MILP Formulation for the Thermal Unit Commitment Problem. *IEEE Transactions on Power Systems*, 28(4):4897–4908, November 2013.
- [7] G. Morales-Espana, A. Ramos, and J. Garcia-Gonzalez. An MIP Formulation for Joint Market-Clearing of Energy and Reserves Based on Ramp Scheduling. *IEEE Transactions on Power Systems*, 29(1):476–488, 2014. ISSN 0885-8950.
- [8] Germán Morales-España. *Unit Commitment: Computational Performance, System Representation and Wind Uncertainty Management*. PhD thesis, Pontifical Comillas University, KTH Royal Institute of Technology, and Delft University of Technology, Spain, September 2014.
- [9] Germán Morales-Espana, Jesus M. Latorre, and Andres Ramos. Tight and Compact MILP Formulation of Start-Up and Shut-Down Ramping in Unit Commitment. *IEEE Transactions on Power Systems*, 28(2):1288–1296, 2013.
- [10] Germán Morales-España, Claudio Gentile, and Andres Ramos. Tight MIP formulations of the power-based unit commitment problem. *OR Spectrum*, pages 1–22, May 2015.
- [11] NERC. "22:00" Frequency Excursions (Final Report). Technical report, 2002.
- [12] NERC. 2014 Long-Term Reliability Assessment. Technical Report November, 2014.
- [13] UCTE. frequency quality investigation, excerpt of the final report. Technical report, 2008.
- [14] Jianhui Wang, Mohammad Shahidehpour, and Zuyi Li. Security-Constrained Unit Commitment With Volatile Wind Power Generation. *IEEE Transactions on Power Systems*, 23(3):1319–1327, August 2008.
- [15] Tobias Weissbach and Ernst Welfonder. Improvement of the performance of scheduled stepwise power programme changes within the European Pow. *IFAC Proceedings Volumes*, 17(1):11972–11977, 2008.