European Journal of Operational Research 254 (2016) 253–268

Contents lists available at ScienceDirect

European Journal of Operational Research

journal homepage: www.elsevier.com/locate/ejor

Innovative Applications of O.R.

Forward thresholds for operation of pumped-storage stations in the real-time energy market

Goran Vojvodic[a,1], Ahmad I. Jarrah[b,*], David P. Morton[c,2]

[a] Department of Decision Sciences, School of Business, The George Washington University, Washington, DC 20052, United States
[b] Department of Decision Sciences, School of Business, The George Washington University, Washington, DC 20052, United States
[c] Department of Industrial Engineering and Management Sciences, Northwestern University, Evanston, IL 60208-3109, United States

ARTICLE INFO

Article history:
Received 12 July 2015
Accepted 13 March 2016
Available online 18 March 2016

Keywords:
Stochastic programming
OR in energy
Large scale optimization
Metaheuristics
Energy markets

ABSTRACT

Pumped-storage hydroelectric plants are very valuable assets on the electric grid and in electric markets as they are able to pump and store water for generation, thus allowing for grid-level storage. Within the realm of short-term energy markets, we present a model for determining forward-looking thresholds for making generation and pumping decisions at such plants. A multistage stochastic programming framework is developed to optimize the thresholds with uncertain system prices over the next three days. Tractability issues are discussed and a novel method based on an implementation of the scatter search algorithm is proposed. Given the size of the multistage stochastic programming formulation, we argue that this novel method is a more accurate representation of the decision process. We demonstrate model stability and quality, and show that the forward thresholds obtained using a stochastic programming framework outperform the forward thresholds from a deterministic model, and thus can lead to efficiency gains for both the generation unit owner and the overall system in the real-time market.

1. Introduction

As capital costs of renewable generation assets continue to decrease and an increasing portion of electrical demand is satisfied by renewable sources, significant attention is being given to storage of electricity. A sizeable share of renewable generation, like wind and solar, is intermittent, and thus those resources become more valuable if one is able to store their output efficiently. Pumped-storage hydroelectric plants are generation assets that can generate power by using water to turn their turbines. They are highly valued on an electric grid because they are also able to pump water and keep it in a reservoir for future use. Hence, they can be used by system operators to balance the load, and can be thought of as huge batteries offering the ability to store energy.

Gains in efficiency at a pumped-storage plant can stem from mechanical and engineering improvements that lead to the plant using less power to pump the water. Financial gains can also be achieved by utilizing the plant more efficiently in day-to-day operations. This study focuses on the latter aspect.

Units that participate in an energy market are most often dispatched based on the "marginal cost" of generating one extra unit of energy. As we discuss in Section 2, this measure of cost is difficult to define for pumped-storage units. As a result, energy price thresholds, which act as a measure of the value of the water, are used by some generation owners to dispatch pumped-storage units. Hence, while thresholds are used in industry, to the best of our knowledge, the current practice is to compute these thresholds using deterministic models and analyses. In this paper, we introduce a stochastic programming framework to compute high-quality thresholds by explicitly considering the stochasticity in the real-time energy prices. We establish the superiority of thresholds determined from the proposed stochastic framework over their deterministically determined counterparts. In the process, we also offer a new approach to obtaining high-quality solutions for large-scale problems of this kind, which are otherwise intractable, using either large-scale decomposition schemes or dynamic programming. Furthermore, we argue that the formulation we solve is a more accurate model of the real-world process because of the way decisions are carried out in practice. We also present an overview of the models and scenario generation methods we developed for the locational marginal prices which we take to be uncertain, and we empirically analyze our modeling framework through stability and quality testing, as well as simulations of real-time operations.

* Corresponding author. Tel.: +1 2029946739.
E-mail addresses: goranv@gwu.edu (G. Vojvodic), jarrah@gwu.edu (A.I. Jarrah), david.morton@northwestern.edu (D.P. Morton).
[1] Tel.: +1 202 994 8520.
[2] Tel.: +1 847 467 2996.

http://dx.doi.org/10.1016/j.ejor.2016.03.020
0377-2217/© 2016 Published by Elsevier B.V.

Table 1
LMPs (/megawatt hour), day-ahead awards and real-time operations (megawatt hour) for unit c and unit p.

Hour (h)	1	2	3	...	Hour (h)	1	2	3	...
(a) Day-ahead LMPs (LMP_h^{DA}) and awards (GEN_h^{DA}).					(b) Real-time LMPs (LMP_h^{RT}) and operations (GEN_h^{RT}).				
Unit c	300	400	350		Unit c	200	200	400	
Unit p	−100	500	200		Unit p	100	0	200	
LMP_h^{DA}	35 dollars	50 dollars	40 dollars		LMP_h^{RT}	50 dollars	30 dollars	40 dollars	

Table 2
Day-ahead and real-time payments for unit c and unit p.

Hour (h)	1	2	3	...	Hour (h)	1	2	3	...
(a) Day-ahead payments.					(b) Real-time payments.				
Unit c	10500 dollars	20000 dollars	14000 dollars		Unit c	−5000 dollars	−6000 dollars	2000 dollars	
Unit p	−3500 dollars	25000 dollars	8000 dollars		Unit p	10000 dollars	−15000 dollars	0 dollar	

We describe the problem and the business case in more detail in Section 2. In Section 3, we give a literature review of related energy applications. We propose our modeling methodology in Section 4, and our algorithm to solve the problem in Section 5. In Section 6, we describe the process used to model energy prices, generate scenarios, and offer evidence of stability and quality. Section 7 deals with the computational performance of the algorithm and discusses the impact the proposed thresholds are expected to have on the operations of the plant. Finally, we conclude with closing remarks in Section 8.

2. The pumped storage real-time decision problem

2.1. Energy market concepts and example

The assumed market environment is referred to as “two-settlement” as it has a day-ahead (DA) and a real-time (RT) component, and each part is cleared, or settled, separately. The market operator is an independent entity responsible for clearing wholesale transactions on the market, such as PJM who operates the world’s largest wholesale market. In the day-ahead market, cost curves are submitted by generation owners, and an hourly generation schedule is awarded by the market operator for the next-day operations. The real-time market is the “day of operations” market. Generation owners are paid based on the deviation of their units’ generation from the day-ahead schedule. Specifically, any real-time generation shortage (excess) from the day-ahead award has to be bought from (sold to) the market.

The dispatching (or operating at a certain output level) of generation units is primarily based on the incremental or “marginal” cost of generating an additional megawatt hour of electric power. Our focus in this research is on determining a coherent dispatching cost for pumped-storage stations for use in the real-time market. Accurately assessing a generation unit’s marginal cost is essential for its profitability and the efficient operation of the market as a whole. Throughout this paper, we assume the electricity market setup is similar to PJM, in which units are dispatched in a least-cost order.

If the estimated cost of generation at power plants is not their true cost, generation units could come online “out-of-merit.” If a particular generation unit is underpriced, its place in the bid stack would be lower than it should be, meaning that it would generate electricity when cheaper resources are not doing so. In the same vein, if a generation unit is overpriced, it would be bypassed in favor of more expensive units. In both of these cases, a waste of resources occurs, and the owner of that generation unit is hurt financially, while the market is not as efficient and cost-effective as it could be.

We explain settlements of the day-ahead and real-time energy markets using a simplified example. In the day-ahead case, the hourly award for a particular generation unit is multiplied by the hourly day-ahead locational marginal price (LMP) to arrive at the payment for that unit. In the real-time case, the payment for a specific unit is expressed as the product of the hourly deviation in output of that unit from the day-ahead award and the hourly real-time LMP.

Assume that unit c is a conventional thermal generator with a marginal cost of generation of 35 dollars per megawatt hour, and unit p is a pumped-storage plant. Assume day-ahead awards for unit c and unit p are given in Table 1a. Positive numbers for awards indicate that energy is supposed to be generated and sold to the market, and negative ones mean that energy is expected to be bought from the market for pumping. Suppose that in real-time, the two plants have the outputs shown in Table 1b. The differences may arise because the plants’ generation is limited by unexpected technical issues or, more commonly, because of the generators’ exploitation of the real-time LMPs to capture more revenue.

The day-ahead payments, for hour h, are given in Table 2a, and are calculated as $LMP_h^{DA} \cdot GEN_h^{DA}$. The real-time payments in Table 2b, are calculated as $LMP_h^{RT} \cdot \left(GEN_h^{RT} - GEN_h^{DA}\right)$. Positive numbers indicate payment to, and negative numbers payment from, the generation unit’s owner. These reflect only revenue, or the “exogenous” component of the profit. From the generation owner’s standpoint, there is also the marginal cost per megawatt generated, or the “endogenous” component of the profit.

Taking both components into account, we can see that the decision to generate 200 megawatt hours under the day-ahead award in hour 2 for unit c was actually a good one even though the generation owner gets penalized 6000 dollars in the real-time market. Had it generated at its day-ahead award of 400 megawatt hours, it would not have made or lost any money in real-time. This would have cost the generation owner 14000 dollars (400 megawatt hours times its cost of 35 dollars/megawatt hour), resulting in a net cost of 14000 dollars for the hour. By generating only 200 megawatt hours, unit c gets penalized 6000 dollars, but internally its cost is only 7000 dollars (200 megawatt hours times its cost of 35 dollars/megawatt hours), which results in a net cost of 13000 dollars for the hour. This example illustrates the benefit of knowing the marginal cost of generation for real-time operations. While that is known for the conventional unit c, it needs to be estimated for the pumped-storage unit p.

2.2. Forward price thresholds

From the real-time perspective, the main reference point is the day-ahead award. The real-time decision problem then becomes whether to deviate from day-ahead awards and, if so, by how