European Journal of Operational Research 255 (2016) 243–258

Contents lists available at ScienceDirect

European Journal of Operational Research

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

Innovative Applications of O.R.

Spatio-temporal hydro forecasting of multireservoir inflows for hydro-thermal scheduling

Timo Lohmann[a], Amanda S. Hering[b], Steffen Rebennack[a,*]

[a] Division of Economics and Business, Colorado School of Mines, Golden, Colorado, USA
[b] Department of Applied Mathematics and Statistics, Colorado School of Mines, Golden, Colorado, USA

ARTICLE INFO

Article history:
Received 24 February 2015
Accepted 6 May 2016
Available online 12 May 2016

Keywords:
Space-time models
Forecasting
Stochastic dual dynamic programming (SDDP)
Stochastic hydro-thermal scheduling
Periodic autoregressive (PAR) model

ABSTRACT

Hydro-thermal scheduling is the problem of finding an optimal dispatch of power plants in a system containing both hydro and thermal plants. Since hydro plants are able to store water over long time periods, and since future inflows are uncertain due to precipitation, the resulting multi-stage stochastic optimization problem becomes challenging to solve. Several solution methods have been developed over the past few decades to compute practically useful operation policies. One of these methods is stochastic dual dynamic programming (SDDP). SDDP poses strong restrictions on the forecasting method generating the necessary inflow scenarios. In this context, the current state-of-the-art in forecasting are periodic autoregressive (PAR) models. We present a new forecasting model for hydro inflows that incorporates spatial information, *i.e.*, inflow information from neighboring reservoirs of the system, and that also satisfies the restrictions posed by SDDP. We benchmark our model against a PAR model that is similar to the one currently used in Brazil. Three multi-reservoir basins in Brazil serve as a case study for the comparison. We show that our approach outperforms the benchmark PAR model and present the root mean squared error (RMSE) as well as the seasonally-adjusted coefficient of efficiency (SACE) for each reservoir modeled. The overall decrease in RMSE is 8.29 percent using our approach for one month-ahead forecasts. The decrease in RMSE is achieved without additional data collection while only adding 11.8 percent more state variables for the SDDP algorithm.

1. Introduction

The optimization of power systems containing both hydro and thermal plants in a regulated market environment, also known as hydro-thermal scheduling, is a well-studied problem (Pereira & Pinto, 1991; Yeh, 1985). In its basic form, a central dispatcher operates the power system and satisfies consumers' electricity demand while minimizing operating costs. The planning horizon of a mid-term hydro-thermal scheduling problem typically spans three to five years with a monthly time resolution and depends on the cumulative storage capacity of the hydro reservoirs in the system relative to total demand. The inflows into the reservoirs of the hydro plants of the system are not known and, therefore, are modeled as uncertain in order to hedge against extreme scenarios. Using too much water for power generation today may lead to high operating costs or, even worse, electricity rationing, in an unforeseen dry period in the future. Whereas, keeping too much water in the reservoirs today may lead to spillage in an unforeseen wet period in the future, again resulting in high operating costs. Similar concepts apply to the case of a deregulated electricity market environment with hydro-power plants in which individual power companies maximize their individual profit. In both cases, the resulting problem can be formulated as a large-scale multi-stage stochastic linear programming (SLP) problem.

Several solution methods have been developed over the past two decades, including scenario tree and various decomposition approaches. In particular, the stochastic dual dynamic programming (SDDP) method (Pereira & Pinto, 1991) is a well-established solution method in this context. The Brazilian power system, for instance, is operated with the aid of a version of SDDP, *cf.* Section 4. The SDDP method is capable of handling a large hydro system, multiple time periods, and an extremely detailed representation of the uncertainty in inflows (Maceira, Duarte, Penna, Moraes, & Melo, 2008). However, it poses the strong restriction on the forecasting method that if inflows are modeled as interstage dependent, they have to be sampled from a linear, additive model (Infanger & Morton, 1996; Queiroz & Morton, 2013). Consequently,

* Corresponding author. Tel.: +1 3032733925.
E-mail address: srebenna@mines.edu (S. Rebennack).
URL: http://www.rebennack.net. (S. Rebennack)

http://dx.doi.org/10.1016/j.ejor.2016.05.011

periodic autoregressive (PAR) type models are commonly used to forecast monthly inflows for a horizon of three to five years ahead (Maceira & Damázio, 2006). Any improvements in the forecasting model will translate into superior optimization results. To date, spatial information has not been directly included in hydro-inflow forecasting models for use with SDDP. We propose and develop such a model; evaluate its performance using real data; and demonstrate how to incorporate such forecasts in the SDDP algorithm.

First, we propose a spatio-temporal approach to forecasting the hydro inflows and evaluate it using three basins from the Brazilian hydro system as a case study. More specifically, instead of only fitting a model to each location's inflows individually, we also use inflow series of neighbors to improve the forecasts at a given location. Selecting neighbors based on their upstream relationship along the river provides the most improvement. The Brazilian hydro system is among the largest in the world with about 140 hydro plants. Historic inflow data is publicly available through the system operator's website (Operador Nacional do Sistema Elétrico, 2015).

We compare our spatial periodic autoregressive (SPAR) model to a benchmark PAR model that is similar to the one used in Brazil (Maceira & Damázio, 2006; Maceira et al., 2008). The main difference between our benchmark PAR and the Brazilian PAR is that the latter is based on an equivalent reservoir scheme that aggregates reservoirs based on geographical location. In order to better compare the PAR to our SPAR, we apply the PAR on the reservoir level instead. We show that our SPAR is superior in terms of more accurate point forecasts while adding little complexity to the SDDP algorithm, *i.e.*, the number of state variables slightly increases or stays the same. Furthermore, our approach does not require additional data collection. For the mathematical programming practitioner who is familiar with hydro-thermal scheduling and the SDDP method, but not necessarily with forecasting inflow scenarios for it, we present a detailed description of our point forecasting approach as applied to a historical inflow data set and provide algorithms for generating forecast scenarios.

Second, we describe how the forecasted inflow scenarios are incorporated into the SDDP algorithm. We derive Benders optimality cuts for our spatio-temporal model and compare them to cuts associated with a PAR model.

The unique contributions of this paper are three-fold.

- First, we describe how the widely used PAR model in Brazil can be made more efficient without even considering spatial information. In particular, we discuss parameter reduction with regard to deseasonalization and model fitting.
- Second, we develop a spatio-temporal hydro-inflow forecasting model for a system of power plants with a monthly fidelity.
- Third, we show how inflow scenarios can be obtained from our model and that they can be incorporated into the SDDP method with little additional effort.

The remainder of this paper is organized as follows. Section 2 reviews the literature on hydro-thermal scheduling and hydro inflow forecasting. Section 3 presents the space-time model along with a description of how to incorporate it into an SDDP routine. Following this, we demonstrate and evaluate the model in Section 4 using the Brazilian hydro system as a case study. Finally, Section 5 offers concluding remarks.

2. Literature review

The main goal of this paper is to introduce and demonstrate the performance of a new forecasting model for hydro inflows that incorporates spatial information. Ultimately, these forecasts would be used in a hydro-thermal scheduling model. Thus, we review the relevant literature in both hydro forecasting and hydro-thermal scheduling in order to motivate our hydro forecasting model choices.

2.1. Solution methods for hydro-thermal scheduling

Several approaches can be found in the literature on how to tackle mid-term to long-term hydro-thermal scheduling problems. Until the 1990s, the predominant method to solve reservoir management problems was the stochastic dynamic programming (SDP) algorithm, which essentially decomposes the multi-stage decision process into decisions for each stage following the principle of Bellman (1957). Different versions of the SDP algorithm emerged, each aimed at mitigating its computational drawbacks (Grygier & Stedinger, 1985; Pereira & Pinto, 1985). Survey papers (Yakowitz, 1982; Yeh, 1985) present a review of hydro-thermal scheduling algorithms in this period. Then, the stochastic dual dynamic programming (SDDP) algorithm was proposed in Pereira and Pinto (1991), which successfully overcomes the biggest drawback of the SDP related methods (at least computationally): the curse-of-dimensionality. The curse-of-dimensionality is the combinatorial "explosion" of the state or scenario space with an increase in (i) reservoirs and reservoir level discretization, (ii) scenarios, and (iii) time stages. The SDDP algorithm has been extensively analyzed, especially with regard to statistical and convergence properties (Philpott & Guan, 2008; Shapiro, 2011).

Extensions on many fronts have been made for the SDDP method. On the inflow scenario front, cut sharing methods were developed to overcome the restriction of independent inflows and to allow interstage dependency (Infanger & Morton, 1996; Queiroz & Morton, 2013). Discrete inflow residuals for each stage have also been proposed to overcome the discretization error exposed to by sampling from a continuous distribution (Pritchard, 2015). The SDDP algorithm has been applied to answer irrigation and environmental questions (Tilmant & Kelman, 2007) as well as to model emission caps (Rebennack, 2014; Rebennack, Flach, Pereira, & Pardalos, 2012). Furthermore, the SDDP algorithm, designed for a regulated market environment, has also been adopted to a deregulated electricity market, wherein instead of a single regulator operating all plants in the system, multiple players (electricity companies) own different parts of the system and submit bids into an electricity wholesale market. In this setting, prices of electricity are typically assumed to be uncertain, giving rise to hybrid SDP–SDDP methods (Gjelsvik, Mo, & Haugstad, 2010; Gjelsvik & Wallace, 1996), bidding strategy analyses (Gjelsvik, Belsnes, & Haugstad, 1999; Scott & Read, 1996; Steeger, Barroso, & Rebennack, 2014; Steeger & Rebennack, 2015), and incorporation of risk aversion Philpott and de Matos (2012); Shapiro, Tekaya, da Costa, and Soares (2013). SDDP can also be applied to nonlinear models such as those that arise when incorporating nonlinear water head effects (Cerisola, Latorre, & Ramos, 2012).

All of the above SDDP-type algorithms share the same restriction with respect to uncertainty in the inflows. If the uncertain inflows are to be modeled as interstage dependent, they must come from a linear additive forecasting model; our proposed model satisfies this requirement. Incorporating improved forecasts in an optimization problem, such as a power system economic dispatch model, has been shown to result in considerable cost savings (Zhu, Genton, Gu, & Xie, 2014b).

SDDP is not the only approach to solve hydro-scheduling problems. In particular, most other methods do not require the hydro inflows to come from a linear additive model. One such method is the scenario tree based method (Boomsma, Juul, & Fleten, 2014; Helseth, Mo, & Warland, 2010), which approximates the random space by a set of scenarios in order to construct a deterministic equivalent and, thus, can use any forecasting model. In addition, constructive dual dynamic programming