Proceeding Series of the Brazilian Society of Computational and Applied Mathematics, v. 11, n. 1, 2025.

Trabalho apresentado no XLIII CNMAC, Centro de Convenções do Armação Resort - Porto de Galinhas - PE, 2024

Proceeding Series of the Brazilian Society of Computational and Applied Mathematics

Boundary Conditions for Hydrothermal Operation Planning Problems: The Infinite Horizon Approach

Bernardo F. P. Costa[1]
IM–UFRJ & FGV EMAp, Rio de Janeiro, RJ
Anderson O. Calixto, Ruan Felipe S. Sousa[2]
IM–UFRJ, Rio de Janeiro, RJ
Ricardo T. Figueiredo[3]
IM–UFRJ, Rio de Janeiro, RJ & UFSC, Florianópolis, SC
Debora D. J. Penna, Lucas S. Khenayfis, Alessandra M. R. Oliveira[4]
ONS, Rio de Janeiro, RJ

Abstract. Multistage stochastic optimization models are used in the energy sector for a variety of planning and operational decisions. Traditional models optimize over a finite planning horizon, which requires setting up ad-hoc terminal conditions, in many cases given as a terminal cost function which induces the water value along the planning horizon. We propose the use of periodic, infinite horizon models, which are better approximations of the long-term behaviour of the system, as a framework to calculate such terminal cost functions. The periodic nature of the model allows for a more efficient algorithm, and leads to a method that is naturally adapted to the system configuration. This is especially useful for the operation planning of hydrothermal power systems, where the water value is a key component of the operation strategy, as well as for the expansion planning of the system.

Keywords. Stochastic Optimization, Value Functions, Periodic Models

1 Introduction

In the energy sector, one of the main applications of stochastic optimization is in the planning and operation of power systems, where the uncertainty in the demand and the availability of renewable energy sources is a major concern. In Brazil, one of the main challenges in the operation of the power system is the large share of hydroelectric generation, which is highly dependent on the availability of water in the reservoirs [10]. Moreover, the increasing share of wind and solar generation, which are highly variable and uncertain, albeit on a different time scale as compared to the hydro inflows, adds to the complexity of the operation planning of the system [7].

In a hydrothermal power system, the operation planning is typically modeled as a stochastic optimization problem, where the objective is to minimize the expected cost of operation, subject to a set of constraints that represent the physical and operational limitations of the system. The operation of the system is coordinated by the National Operator (ONS), which is responsible for the dispatch of around 160 hydroelectric and 130 thermal power plants, the management of the transmission grid and taking into account wind, solar and distributed generation [6].

Due to the size of the system and the complexity of the problem, the operation planning of the power system is divided into a chain of models, which are solved sequentially, each one with a

[1]bernardo.paulo@fgv.br
[2]acalixtorcp@gmail.com, ruanfelipe@matematica.ufrj.br
[3]ricardoturanofigueiredo@gmail.com
[4]debora.jardim@ons.org.br, lkhenayfis@ons.org.br, amro@ons.org.br

different time horizon and level of detail [8]. The main models in use today for operation planning **start** from a 5-year stochastic optimization model, which incorporates the uncertainty in the inflows and availability of the power plants, and was chosen due to the multi-year regulation capacity of the hydro reservoirs and dry/wet cycles. This first model yields a set of **Future Cost Functions**, that represent, for each month of the planning horizon, the expected cost of operation given the energy stored in the reservoirs and the hydrological scenario. The Future Cost Function at the end of the second month is then used as input for the **second main model** in the aforementioned chain, a 2-month model, which has a higher level of detail of the hydro plants, besides having weekly time steps for the first month. This 2-month model then produces its own Future Cost Functions, at the end of each week of the first month, which is then used as input for the **third** operation planning model of the chain, a 1-week model, with hourly timesteps in the first day and daily timesteps until the end of the week, which is used to define the operation strategy for the next day.

We focus on the first model in the chain, since its results impact the other models, and in particular its Future Cost Functions, which are used as input for the second model. The current finite horizon model needs a careful treatment of the terminal conditions, given as a **terminal cost function**, whose choice needs to be justified. The simple choice of a zero function would imply that the water stored in the reservoirs at the end of the horizon has no value, which leads to an "end-of-world" effect, where the system would not refrain to empty the reservoirs at the end of the horizon, which is not realistic. Therefore, in Brazil's current practice, after the 5-year planning horizon, another 5 years are added, with a fixed system configuration, and the terminal condition for the resulting 10-year problem is set to zero. This effectively corresponds to pushing the "end-of-horizon" effect to the end of the 10-year horizon, what is believed to be enough to make its effect negligible for the first 5-years of interest.

Indeed, in principle, a sufficiently long planning horizon would allow for the first stages of the system to become almost independent of the ad-hoc terminal cost function. However, the determination of such appropriate horizon depends (at least) on the system configuration and the magnitude of the uncertainties, and therefore the choice of the extra 5 years needs to be justified. Given that the present system is already different from when the additional 5-year approach was established, and that the system is expected to continue to change in the near future, we anticipate the need for a more reliable method for setting its terminal condition.

Contributions

We propose the use of an infinite horizon model to setup the boundary conditions of the long-term operation planning of hydrothermal systems in a principled way. Naturally, a different model also impacts the algorithm used to solve the corresponding optimization problem. Fortunately, the standard dynamic programming approach to solve stochastic optimization models can be adapted to an infinite horizon model, by using a **discount factor** to account for the future costs [1]. This discount factor is already used in the finite horizon models, and corresponds to an annual interest rate. Moreover, the periodic nature of the hydrothermal system, with the annual cycle of inflows and demand, allows for a **finite** representation of the infinite horizon model.

It is the use of a finite and periodic problem that is at the core of our proposal. This has several advantages over the current practice of using a finite horizon model. Regarding the value functions:

- A simple, 1-periodic, example illustrates the difference between the value functions, and especially in the water values;
- A more complex 12-periodic problem in higher dimension highlights the benefits of a more coherent treatment of the water value in the operation planning, leading to a more balanced

thermal dispatch and reduced load shedding.

Furthermore, the optimization of such periodic models can typically be performed in a separate step relative to the operations planning, therefore reducing the time between the availability of new data and the decision making.

2 Models

The mathematical model for the operation planning of a hydrothermal system is a multistage stochastic optimization problem, where the state variables represent the energy stored in the reservoirs, and the control variables represent the dispatch of the generation units and transmission between network nodes. The uncertainty is represented by a set of scenarios, which are assumed to be independent between each stage in order to use dynamic programming with a single value function for each stage. The objective is to minimize the expected cost of operation, subject to a set of constraints that represent the physical and operational limitations of the system.

In the finite horizon setting, this is represented by the following set of equations:

$$V_{T+1}(x) = 0, \tag{1}$$

$$V_t(x,\omega) = \min_{u\in\mathcal{U}} \left\{ c_t(x,u,\omega) + \beta V_{t+1}(f_t(x,u,\omega)) \right\}, \tag{2}$$

$$V_t(x) = \sum_{\omega\in\Omega} p_t(\omega) V_t(x,\omega), \qquad \forall t = 1,\dots,T, \tag{3}$$

where $V_t(x)$ is the value function at stage t and state x, u is the control variable, $c_t(x,u,\omega)$ is the cost of operation, $f_t(x,u,\omega)$ is the transition function, and $p_t(\omega)$ is the probability of scenario ω at stage t, so that $V_t(x)$ is the average of the cost-to-go for each scenario $V_t(x,\omega)$. The parameter $0 < \beta < 1$ is the discount factor, used to account for the future costs at present value.

In the Brazilian practice, $T = 120$ stages, corresponding to 10 years of monthly stages, with 5 years of planning and 5 extra years whose purpose is to yield $V_{61}(\cdot)$ as the terminal condition for the 5-year model. Moreover, costs and constraints are represented as a Linear Programming model, with the transition function given by a linear equality constraint. This is a simplification which is justified by good linear approximations of the dynamics, and moreover by the mathematical properties of the value functions V_t, which, in this setting, are **convex** in the state variable.

An infinite horizon model has period τ when the costs, constraints, and transition function are τ-periodic, that is, $c_t(x,u,\omega) = c_{t+\tau}(x,u,\omega)$, and similarly for the other functions. If the states and controls are bounded, then the expected cost over the entire horizon is also bounded, because $|\beta| < 1$. In this case, the value functions are also τ-periodic, that is, $V_t(x) = V_{t+\tau}(x)$. Therefore, the dynamic programming recursion of equations (1)–(3) are replaced by the system of equations

$$V_t(x,\omega) = \min_{u\in\mathcal{U}} \left\{ c_t(x,u,\omega) + \beta V_{t+1}(f_t(x,u,\omega)) \right\}, \tag{4}$$

$$V_t(x) = \sum_{\omega\in\Omega} p_t(\omega) V_t(x,\omega), \qquad \forall t = 1,\dots,\tau, \tag{5}$$

which is a fixed-point problem since $V_{\tau+1} = V_1$, which makes the equation (4) for $t = \tau$ depend on V_1, contrary to the finite horizon model.

The standard algorithm for solving the stochastic finite horizon model, the SDDP algorithm [9], is based on dynamic programming and lower approximations of the value function. It can be adapted to an infinite horizon model, replacing the dynamic programming recursion with a fixed-point equation [2]. Mathematically, the solution of the finite-horizon model starts from the terminal value function and recursively computes the value function at each time step, while the solution for the infinite-horizon model starts from an initial guess of the value function and iteratively applies a fixed-point operator, which is a contraction mapping, and converges to its value functions. While this looks like a different approach, the algorithmic implementation for both models is very similar, starting from a valid lower bound at each stage, and iteratively improving it by adding cuts [4]. The only difference is that, in the periodic setting, since the value functions are T-periodic, a cut for stage t can be used for all stages $t + kT$. In particular, we need only represent a single period of the problem. Further details, references and variants can be found in [3].

3 Numerical Results

We present two different models to analyse the impact of the periodic setting. First, a toy model, where each stage represents a year, with a single state variable; then a more realistic model, where each stage represents a month, and has a 4-dimensional state space. The toy model was obtained from the monthly model by aggregating the entire system into a single reservoir, and considering yearly timesteps, and was considered as a sanity test.

For simplicity, both models have periodic data, which means that we do not consider the expansion of the system which is usually incorporated in the Brazilian practice. Since our main objective is to compare the periodic model with a finite horizon one, this does not impact our analysis: For the value functions at the end of the **planning horizon**, both problems will have periodic data, since the 5-year extension does not change the system configuration.

3.1 Toy model

A periodic model and a 10-year model were both solved using 400 iterations of the SDDP algorithm, and a discount factor of $\beta = 0.9$, using the Julia library `SDDP.jl` [5]. The number of iterations is sufficiently large for such a simple problem: both models improved by less than 10^{-3} in the last 50 iterations.

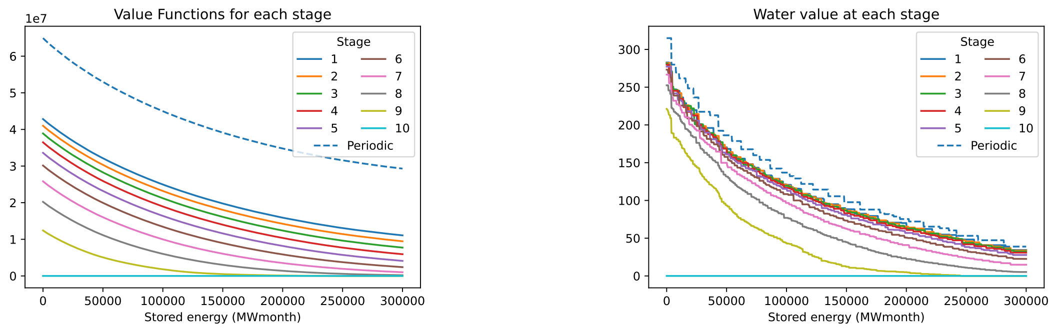

(a) Value functions. (b) Water values.

Figure 1: different stages for the 10-year model, and water value for the periodic model.

As we can see from Figure 1, both the value functions V_t and the water value (equal to $-\partial_x V_t$) in the finite horizon model increase as we go back in time, as expected from moving away from the end of horizon with a zero boundary condition. The increase of the value functions is explained by the fact that the dynamic programming recursion is a contraction mapping, and therefore V_{10}, V_9, $\dots$, V_1 are the ten first elements of a sequence converging to a positive fixed point, moving away from zero. The water value, however, after a rapid start, stagnates around stage 5, which could be taken as a sign that the 5-year extension is sufficient to capture the long-term behaviour of the system. Indeed, the only impact of the value function on the decision of the system is through its derivative. However, as we also observe from Figure 1b, the water values at stages 1 through 5 remains consistently lower than the water value in the periodic model. As a consequence, the value function at the 5-th stage is still affected by the terminal condition at year 10, and moreover the decision at the very first year is also impacted by a lower water value.

3.2 4d model

In the monthly model, we consider a 4-dimensional state space, which represents the energy stored in each of the 4 **equivalent reservoirs**, aggregating basins, respectively, in the SE, S, NE and N regions of Brazil. Due to the dimension of the value functions, we show the average operation of the system using both policies, from a set of 10 000 scenarios, to assess the impact of switching to the periodic model. In Figure 2, we present the results of the 120-stage model and its 12-periodic counterpart, which were both solved using 500 iterations of the SDDP algorithm, with an annual discount factor of $\beta = 0.9$.

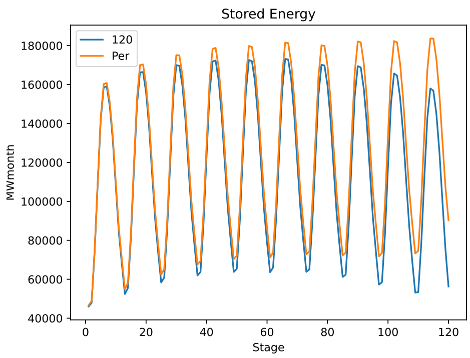

(a) Average stored energy at each stage.

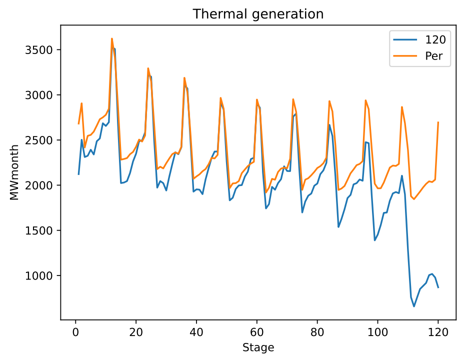

(b) Average thermal generation at each stage.

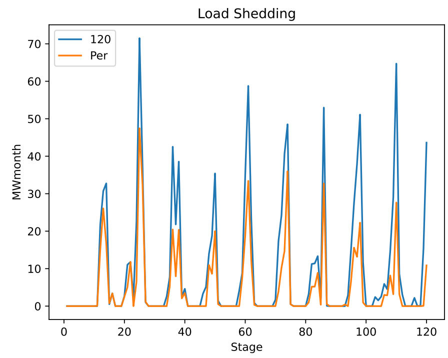

(c) Average load shedding at each stage.

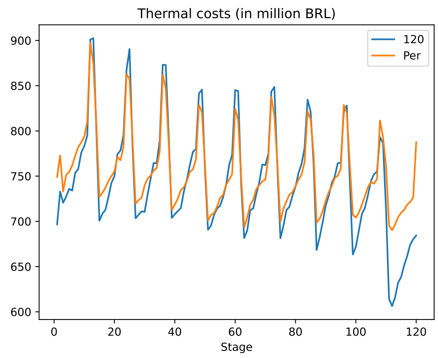

(d) Average generation costs at each stage.

Figure 2: Simulation results for the 120-stage model and its 12-periodic counterpart.

As we can see in Figure 2a, the average stored energy is higher in the periodic model, which is a consequence of the higher water value. This is coupled with a higher thermal generation, as displayed in Figure 2b, especially leading to higher **preventive thermal generation** during the

DOI: 10.5540/03.2025.011.01.0355

wet seasons. As a consequence, even if the peak average thermal generation is not significantly different in both settings, the finite horizon model incurs significantly more load shedding, as we can see in Figure 2c. Finally, it is also important to notice that the peak generation costs of Figure 2d are also higher in the finite horizon model. This is due to the non-linear nature of these costs, as thermal units are generally dispatched by increasing unit costs; therefore, a more balanced thermal generation from the periodic model results in a lower average cost.

We also observe that the system starts from a modest storage level, which has a tendency to increase in the first years of the horizon, and then stabilize around a higher level. The level where this stabilization occurs is higher in the periodic model, which is a consequence of the higher water value and thermal generation. Again, if one had only access to the finite horizon model, one could be tempted to think that the relatively stable storage levels around years 4–6 represent the long-term behaviour of the system, but the region where it stabilizes is still affected by the terminal condition at year 10. This is very clearly seen in Figure 3a, where we moreover observe that the average stored energy along each year still shows an increasing behaviour between years 5 and 10. This indication that 5 years is not enough to reach a steady state, even in the periodic model, is further confirmed by calculating the distance between the cumulative distribution functions of the stored energy at the end of the wet season. As we see in Figure 3b, these distances only start to become reasonably small after year 5, and still year 6 is further away from the steady state than the following years.

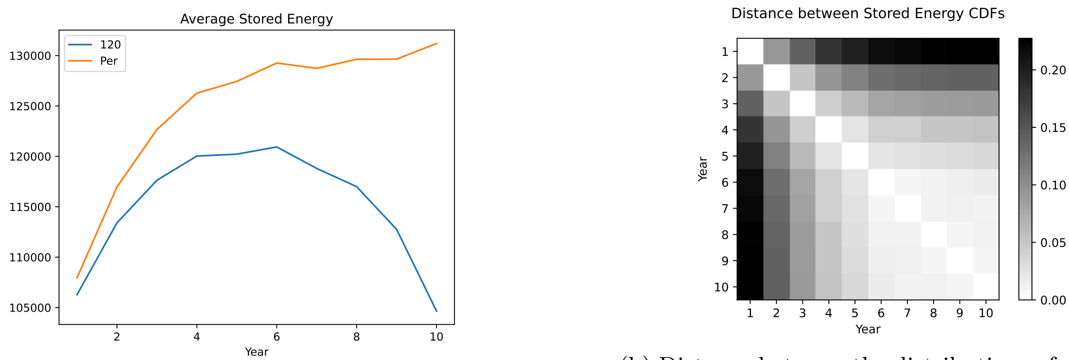

(a) Simulation of the stored energy, averaged for each year, for both models.

(b) Distance between the distributions of stored energy at the end of the wet season, for the periodic model

Figure 3: Long-term behavior of the stored energy for the 120-stage and the 12-periodic infinite-horizon models.

4 Concluding remarks

We proposed to setup the terminal condition for the long-term planning of hydrothermal systems through an infinite horizon model. Besides the intuitive advantage of definitely removing the end of horizon, this model is also parameter-free, inasmuch as it does not depend on the system configuration. Also, we observed that the finite horizon model has results which are still impacted by the terminal condition, even after a 5-year extension, which was only possible because we used the infinite horizon model to establish a sound basis for comparison. Moreover, if only the finite horizon model were used, one could misinterpret the observed stability of the system around year 5 as an indication of an equilibrium, which the periodic model shows to be, first, too low in terms of water storage, and second, too early to be reached.

Furthermore, the time-invariant nature of the periodic model is also more useful for expansion planning, as it captures the intrinsic nature of the system, and not a specific realization of the inflows. In particular, the periodic value functions are independent of the initial condition of the system; and the state variables follow a Markov chain under the optimal policy, so its "steady-state" captures the long-term behaviour of the variables. We expect that these results could be explored in future work.

Acknowledgements

Several results were obtained in 2021 and 2022 in the setting of project IM-23145 between UFRJ and ONS, and the authors would like to thank the support for developping these themes. Earlier versions of this work were presented at a technical meeting of CPAMP in October 2023, and at CEPEL in March 2024, for which we thank the inspiring discussions. The authors would also like to thank Joari Paulo da Costa for numerous discussions during this work.

References

[1] D. P. Bertsekas. **Dynamic programming and optimal control: Approximate dynamic programming**. 4th. Athena Scientific, 2012.

[2] J. Birge and G. Zhao. "Successive Linear Approximation Solution of Infinite-Horizon Dynamic Stochastic Programs". In: **SIAM Journal on Optimization** 18.4 (Jan. 2007), pp. 1165–1186. DOI: 10.1137/060665506.

[3] B. F. P. Costa, A. O. Calixto, R. T. Figueiredo, R. F. S. Souza, D. D. J. Penna, A. M. R. Oliveira, and L. S. Khenayfis. **Problemas de otimização estocástica em horizonte periódico**. Relatório Técnico. Convênio ONS – UFRJ, 2022.

[4] O. Dowson. "The policy graph decomposition of multistage stochastic programming problems". In: **Networks** 76.1 (Feb. 2020), pp. 3–23. DOI: 10.1002/net.21932.

[5] O. Dowson and L. Kapelevich. "SDDP.jl: a Julia package for Stochastic Dual Dynamic Programming". In: **Optimization Online** (2017). URL: http://www.optimization-online.org/DB_HTML/2017/12/6388.html.

[6] ONS. **Atuação do ONS**. Online. Acessado em 18/03/2024, https://www.ons.org.br/paginas/sobre-o-ons/atuacao.

[7] ONS. **O Sistema em Números**. Online. Acessado em 18/03/2024, https://www.ons.org.br/paginas/sobre-o-sin/o-sistema-em-numeros.

[8] M. V. F. Pereira. "Optimal Scheduling of Hydrothermal Systems - An Overview". In: **IFAC Proceedings Volumes** 18.7 (1985). IFAC Symposium on Planning and Operation of Electric Energy Systems., Rio de Janeiro, Brazil, 22-25 July, pp. 1–9. ISSN: 1474-6670.

[9] M. V. F. Pereira and L. M. V. G. Pinto. "Multi-stage stochastic optimization applied to energy planning". In: **Mathematical programming** 52.1-3 (1991), pp. 359–375.

[10] L. A. Terry, M. V. F. Pereira, T. A. Araripe Neto, L. F. C. A. Silva, and P. R. H. Sales. "Coordinating the Energy Generation of the Brazilian National Hydrothermal Electrical Generating System". In: **Interfaces** 16.1 (1986), pp. 16–38.

