Technical Session 5A (15:20 - 16:30)	**Enhancements in Long Term planning and SDDP**	**Chair: Vitor de Matos, NORUS**	
	Enhanced scalability in SDDP-based hydropower scheduling: A performance analysis on high-performance computing infrastruc-ture	Knut Gjerden	SINTEF
	Novel Cut Selection Techniques for Stochastic Dual Dynamic Programming using Pareto Dominance Arguments	Janik Königshofer	KIT
	Strategies for an Efficient Representation of the Piecewise Linear Approximation of the Hydropower Production Function in Long-Term Operation Planning	Pedro Lira	CEPEL

Enhanced scalability in SDDP-based hydropower scheduling: A performance analysis on high-performance computing infrastructure

Knut Skogstrand Gjerden,[1]

[a]SINTEF Energy Research, Department of Energy Systems, Postboks 4761 Torgarden, Trondheim, 7465, Norway

---

Abstract

We present significant advancements in the parallel performance of an SDDP-based algorithm allowing us to leverage the increased computational power offered by modern high-performance computing (HPC) infrastructure. The optimization involved in utilizing these powerful systems includes parallel work partitioning strategy and parallel task distribution. This work utilizes the Betzy supercomputer, Norway's most powerful supercomputer with a peak theoretical performance of 6.2 petaflops, to achieve substantial improvements in execution speed on a larger number of processes compared to previous implementations. We demonstrate a notable reduction in runtime as well as a larger region of number of processes with sustained parallel performance. The analysis also highlights the scalability of the algorithm on large datasets and suggests its potential for wider applications with improved computational efficiency. This paper provides valuable insights into leveraging HPC resources for accelerating complex algorithms, paving the way for future research and development in this area.

Key words: High-performance computing, Algorithm optimization, Parallelism, Computational efficiency, Scalability.

Email address: `knut.gjerden@sintef.no` (Knut Skogstrand Gjerden)

Novel Cut Selection Techniques for Stochastic Dual Dynamic Programming using Pareto Dominance Arguments

Königshofer, Janik   Warwicker, John   Füllner, Christian   Rebennack, Steffen

Karlsruhe Institute of Technology

Abstract

Sequential decision problems arise in numerous real-world applications, often under conditions of uncertainty. These problems require decisions to be made while accounting for potential future outcomes, such as the operation planning of power plants in the energy industry, where stochastic factors like weather significantly influence renewable energy production. Stochastic Programming provides a powerful framework for modeling such problems, enabling the use of advanced decomposition-based methods to address large-scale, multi-stage models in a computationally efficient manner. Among these, Stochastic Dual Dynamic Programming (SDDP) is regarded as the state-of-the-art algorithm for solving large-scale multi-stage stochastic programs. Efficient solution times are critical for enabling experimentation with diverse future scenarios and incorporating newly available information. While recent efforts to accelerate SDDP through machine learning approaches have shown promise, these methods can undermine the algorithm's theoretical guarantees regarding bound properties and convergence. Cut selection methods offer a complementary approach to speed up SDDP without sacrificing these guarantees. This thesis introduces a novel cut selection framework based on Pareto dominance arguments, offering an alternative to the established Level One cut selection approach. Several instances of this framework are proposed and benchmarked against standard SDDP and Level One cut selection using multiple instances of the Brazilian hydro-thermal scheduling problem with varying problem sizes and initializations. Empirical results demonstrate that Pareto-dominance-based strategies can outperform current methods in specific scenarios, requiring only 80% of the time to reach a certain bound quality compared to the state-of-the-art approach.

Strategies for an Efficient Representation of the Piecewise Linear Approximation of the Hydropower Production Function in Long-Term Operation Planning

Pedro T. M. Lira[1], André Luiz Diniz[1], Lílian C. B. Santos[1], Cristiane B. Cruz[1]
[1]Electric Energy Research Center - CEPEL

I. Introduction

In the operation planning of the Brazilian electric system, which is predominantly hydroelectric, it is essential to accurately model the nonlinear generation of hydroelectric plants ("Exact Hydropower Production Function"), which depends on the net head and the turbine flow rate [1]. Due to its complexity and the challenges involved, this planning is divided into long-, mid-, and short-term stages, each with different levels of detail.

However, due to the characteristics of the official long-term model (NEWAVE) used by the Brazilian ISO to define the dispatch policies and by the Chamber for Commercializing Electricity (CCEE) for price calculation (PLD), the Hydropower Production Function (HPF) representation must be formulated using linear programming, which currently includes two modeling approaches:

i. Piecewise linear approximation: a representation based on the stored volume, turbine flow rate, and spillage for each hydroelectric plant, referred to as the Approximate Hydropower Production Function (AHPF) [1]. This modeling approach is officially used in NEWAVE, as well as in DECOMP (mid-term) and DESSEM (short-term/daily scheduling). Due to its higher level of detail, the AHPF requires significantly greater computational effort. Equation (1) defines the AHPF for each hydro plant i, where α is a factor determined using a linear regression technique, v is the storage, q is the turbined outflow, s is the spillage, γ are the coefficients for each variable and K_i is the number of hyperplanes:

$$gh_i \leq \alpha_i \left(\gamma_{0,i}^k + \gamma_{v,i}^k v_i + \gamma_{q,i}^k q_i\right) + \gamma_{s,i}^k s_i,\ k = 1, \ldots, K_i \quad (1)$$

ii. Constant productivity: a linear function of turbine flow rate that does not capture the variation in net head. This approach is highly simplified and advantageous in terms of computational time but results in significant deviations from the exact HPF. The Equation (2) defines this approach, where ρ_i is the constant productivity of each hydro plant i, calculated using a seasonal reference volume for each month:

$$gh_i = \rho_i q_i \quad (2)$$

The parameters used to construct the AHPF for each plant vary between models, reflecting differences in the planning horizon and the level of detail applied. For instance, in the long-term model, the storage volume window covers the total usable volume of the plant, while in the mid-term model, it is defined within a ±10% range around the initial volume. The turbined outflow window extends from the minimum to the maximum values in both models. Despite these differences, the number of hyperplanes used to represent the AHPF for each plant is quite similar across models, as both employ five points to discretize the domains of usable storage and turbined outflow when constructing the convex hull.

II. Methodology

The piecewise linear representation used to represent the AHPF introduce multiple constraints (one for each hyperplane), increasing the complexity of the problem. Therefore, the motivation of this work is to preserve as much detail as possible in the representation while achieving meaningful reductions in computational effort. To this end, two methodologies are proposed, applied to the long-term planning model and extendable to mid- and short-term models, to reduce the total number of hyperplanes:

- A comparison of the **angles** between the normal vectors of two hyperplanes. If the difference is smaller than a predefined threshold, they are replaced by an average hyperplane.
- A comparison of the generation values at a set of points for each pair of hyperplanes. If the maximum distance between their **projections** exceeds a predefined threshold, the two hyperplanes are merged into a single average hyperplane.

For these two methodologies, we construct the convex hull using fifty points to discretize the domains of usable storage and turbined outflow, allowing for a more accurate representation of the exact function's behavior. Although this procedure initially generates a large number of hyperplanes, most of them are eliminated after applying the proposed methodologies.

Finally, a performance analysis was conducted on the two HPF models, along with the proposed improvements, using the ”hybrid” NEWAVE (one year with individualized representation of hydroelectric plants and four years with energy equivalent reservoirs representation)

III. Results and Conclusions

To compare the results of the proposed methodologies, we conducted an analysis using data provided by the Brazilian ISO in January 2024 with the NEWAVE model, considering 155 hydroelectric plants, 60 periods with monthly discretization, 20 realizations per period, and 200 forward resampling scenarios in the Stochastic Dual Dynamic Programming (SDDP) algorithm. After the convergence of the SDDP algorithm, a final simulation with 2000 scenarios was conducted to evaluate the performance of the methods, using the same AHPF for all cases: the current (reference) model. Figure 1 presents the time per iteration for each methodology during the SDDP process, while the Figure 2 illustrates the empirical cumulative distribution function using average absolute deviations for each plant in the final simulation. Table I shows the total time of the SDDP algorithm and the global average absolute deviations in MW and percentage. These results indicate that, although the constant productivity model achieves approximately 50% lower computational effort compared to the reference model, the average absolute deviations are approximately six times higher, indicating a significant loss of precision when disregarding the effects of variation in net head.

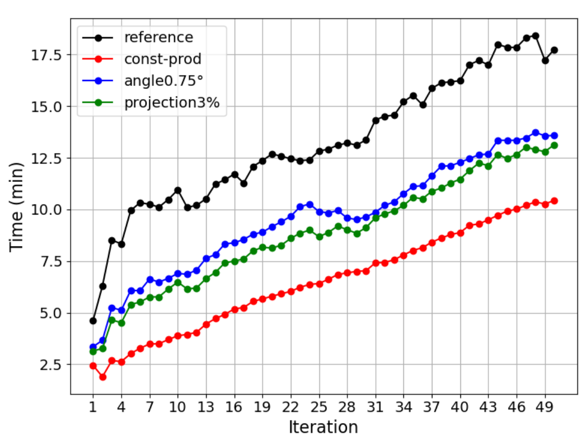

Fig. 1. Time per iteration for each case.

Additionally, the proposed methodologies significantly reduced computational effort, with slight increases in the deviations between the AHPF and the exact HPF. To illustrate the potential of this approach, for hybrid cases using a criterion of an angle smaller than 0.75° between the normal vectors and a projection onto the generation axis lower than 3% of the installed capacity of each hydro plant, computational time was reduced by approximately 28% and 34%, respectively.

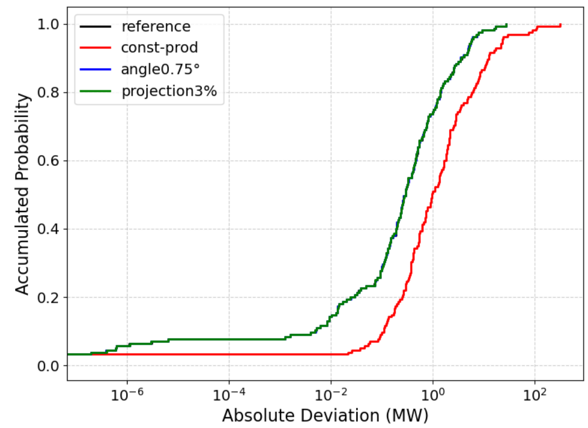

Fig. 2. Empirical cumulative distribution function of the average deviation per plant for each case (x-axis is in logarithmic scale). Black, blue and green graphics are overlapping.

TABLE I
SDDP solution time and global average absolute deviation.

Methodology	Time (min)	Average deviation
reference	662	1.3439 MW 1.275%
constant productivity	325	7.6411 MW 6.993%
angle 0.75°	479	1.3411 MW 1.272%
projection 3%	443	1.3440 MW 1.275%

Meanwhile, the difference in average deviations is neglegible, highlighting the strong cost-benefit of the methodologies. It is important to note that these graphs can “slide” between the current AHPF representation and the constant-productivity HPF depending on the chosen values for each criterion. For example, if an angle of 1.0° were used, the total number of hyperplanes would decrease, reducing both the problem size and the computational effort required to solve it. Consequently, the time-per-iteration curve in Figure 1 would shift downward. On the other hand, the average deviations would likely increase, causing the graph to shift to the right in Figure 2.

Since long-term operation planning is traditionally solved using SDDP algorithm due to the large problem dimension and a planning horizon spanning several years, computational effort is always a critical concern. The results presented illustrate that it is possible to maintain a reasonable level of detail in representing the variation in net head through the AHPF while simultaneously reducing computational time in solving these problems.

References

[1] A. L. Diniz, M. E. P. Maceira. “A Four-Dimensional Model of Hydro Generation for the Short-Term Hydrothermal Dispatch Problem Considering Head and Spillage Effects”, *IEEE transactions on power systems*, v. 23, n. 3, pp. 1298–1308, 2008.