

Applying High-Performance Computing to the European Resource Adequacy Assessment

Daniel Ávila, Anthony Papavasiliou, *Senior Member, IEEE*, Mauricio Junca, and Lazaros Exizidis¹

Abstract—This work considers the European Resource Adequacy Assessment, which is a pan-European resource adequacy process that is being developed by the European Networks of Transmission System Operators for Electricity (ENTSO-E). A critical part of this process is the so-called Economic Viability Assessment model, which aims at determining future expansion and retirement capacity opportunities for the entire European network. As such, the problem is stochastic. Nevertheless, due to computational constraints, simplified approaches have been followed by ENTSO-E. Our work formulates the problem as a two-stage stochastic problem and proposes two decomposition algorithms for solving the problem which are implemented in a high-performance computing infrastructure. The first is a subgradient-based algorithm, and the second uses a relaxation of the second stage (the economic dispatch) in order to speed up the subgradient calculation thus achieving a considerable reduction in solution time. We compare our schemes against the commonly used Bender's decomposition. We compare the obtained stochastic solution against the deterministic solution proposed by ENTSO-E for their 2021 study and analyze the impact of the stochastic solution on various adequacy indicators.

Index Terms—Stochastic capacity generation expansion, Parallel & Decomposition algorithms, Large-scale optimization

I. INTRODUCTION

WITHIN an uncertain world, measuring and analyzing the ability of the electric power system to react to adverse uncertain conditions has become increasingly important. The Clean Energy Package [1] has recognized the importance of this task with Regulation (EU) 941\2019 [2] and Regulation (EU) 943\2019 [3]. The latter stipulates the need for a robust European resource adequacy assessment that provides an instrument for detecting and measuring adequacy concerns [4]. In particular, resource adequacy concerns identified through the European Resource Adequacy Assessment (ERAA) are to become the basis for justifying the implementation of capacity mechanisms within European Member States. As required in Regulation (EU) 943\2019 [3], Member States wishing to introduce capacity mechanisms must do so on the basis of an adequacy concern that is identified in the ERAA study, complemented possibly with a national resource adequacy study. Consequently, there is an institutional urge to develop a reliable and robust ERAA study at a pan-European level.

The European Network of Transmission System Operators for Electricity (ENTSO-E) is the body mandated by regulation to develop the methodology and conduct the study on an annual basis [5]. The ERAA is a pan-European resource

adequacy assessment for up to 10 years ahead, which aims at measuring the ability of the power system to react to a set of future uncertain events [4]. The ERAA study covers the entire pan-European interconnected system, thus modeling 56 bidding zones in 37 countries. Adequacy concerns are identified with the help of adequacy indicators, with the Loss of Load Expectation (LOLE) being the most common indicator used among the EU Member States to define the respective reliability standards. The Expected Energy not Served (EENS) indicator is also computed in the ERAA study, in order to assess the depth (MW) of the curtailments. Such indicators are naturally linked to the installed capacity mix. For the purpose of determining the installed capacity mix, the ERAA study introduces the so-called Economic Viability Assessment (EVA), which aims at modeling economic parameters that affect the available generation capacity within Europe. Unfortunately, due to the scale of the ERAA, incorporating an EVA that integrates the stochastic nature of future events has thus far been considered as being out of scope. Therefore, a deterministic approach has been followed in ERAA 2021 [6], while in ERAA 2022 a tractable formulation is obtained by considering a reduced stochastic model, this simplification consists of using 3 of 35 uncertainty realizations. ENTSO-E is moving towards a stochastic programming formulation [7], endorsed by ACER, but currently limited to three scenarios. This motivates the development of a framework which allows tackling the EVA in its stochastic version.

Our work aims at bridging this methodological gap by proposing a novel parallel computing algorithm, based on ideas from stochastic programming, and implemented on high-performance computing infrastructure. Our approach allows us to account for the stochastic nature of the EVA study in ERAA. Furthermore, we study the potential impact of ignoring the stochastic nature of the EVA on the capacity mix, and in turn, the consequences that this could have on adequacy indicators.

The EVA aims at determining capacity expansion and capacity retirement opportunities for the entire European network. As such, it relates to two streams of literature: (i) stochastic capacity generation expansion, with the added complexity of considering retirement opportunities; (ii) large-scale optimization, which is tackled with the aid of high-performance computing.

A. Stochastic capacity generation expansion

The stochastic capacity generation expansion literature presents a variety of strategies for addressing generation

¹ The information and views set out in this paper are those of the authors solely and do not necessarily reflect the official opinion of ENTSO-E

expansion¹. When the size of the problem is manageable, the problem is solved directly by a commercial solver [8]–[11]. In order to decrease the computational burden, certain authors have followed scenario selection techniques, and report solving time improvements with corresponding losses in the quality of the solution [10], [12]–[17]. Note that these approaches typically still rely entirely on a commercial solver for solving the reduced problem. By relying on stochastic programming formulations [18], certain authors use Bender's decomposition in order to tackle the problem [19]–[22]. Furthermore, modelling tools have been developed in order to solve the problem pragmatically. Examples of open-source software include the EMPIRE model [23] which solves the problem as a large-scale linear problem, while [24] has proposed extensions of EMPIRE where the problem is decomposed by using progressive hedging. Recent approaches [25], applied to the Brazilian power system, have proposed extensions to Benders decomposition in order to speed up convergence. Examples of commercial tools include the PLEXOS software, which allows solving the problem as a large-scale linear problem or using Bender's decomposition. Certain heuristics have also been proposed, including a rolling-horizon scheme [26] and an hourly aggregation of time series [27].

The modelling specifications of the EVA prevent us from solving the problem directly using a commercial solver. These specifications result in a large-scale problem, due to its wide geographical scope, combined with the time step chosen by ENTSO-E, as well as the endogenous representation of uncertainty. We highlight that a small set of uncertainty realizations are already enough to produce a problem of considerable size, thus preventing the use of a scenario reduction technique which allows tackling the problem in its extended form without using a decomposition strategy. On the other hand, a drawback of Bender's decomposition is that its performance diminishes as the number of expansion/retirement candidates increases. The progressive hedging algorithm presented in [24] appears to deteriorate in performance as the size of each scenario subproblem increases, which renders it impractical for our purposes.

Some early theoretical work from 1986 [28] presented the possibility of employing subgradient schemes in deterministic capacity expansion problems. Furthermore, past research has proven the effectiveness of subgradient algorithms when dealing with large instances of stochastic unit-commitment [29], [30]. Consequently, our work aims at translating this success into the ERAA study. Due to the EVA modelling specifications, the calculation of each subgradient is time-consuming. Therefore, our work, in addition to putting forward a subgradient-based algorithmic framework, proposes a further algorithm that calculates approximations of the subgradients efficiently, thus providing notable computational benefits.

B. High performance computing

High-performance computing has become critical in tackling large-scale power system optimization problems. Early

work on the topic [31], [32] describes parallel computing schemes for addressing security-constrained optimal power flow and hydro systems, respectively, using Bender's decomposition. Parallel computing schemes for Lagrangian decomposition have been developed for solving optimal power flow problems [33], [34]. In recent years, parallel computing has enabled tackling large-scale instances of stochastic unit commitment [35]. In particular, the usage of asynchronous parallel computing has resulted in a reduction of computation time from weeks to a few hours for certain stochastic unit commitment instances [30]. The hydrothermal scheduling literature has benefited from the usage of high-performance computing as well. While the known Stochastic Dual Dynamic Programming (SDDP) algorithm [36] has been developed for tackling such a problem, as explained in [37] SDDP can encounter challenges, and several parallel computing studies have followed. For instance, in [38] the authors propose a synchronous parallel scheme, while [39], [40] have proposed asynchronous implementations. In [41] the authors have compared these different parallel implementations. Some recent work [42] has shown further parallelization strategies that are capable of outperforming the commercial implementation of PSR-SDDP² which is targeted for hydrothermal scheduling.

The parallelization of decomposition techniques such as Bender's decomposition or subgradient methods has been studied in the past [31], [32], [34], [35], [43]. Nevertheless, the stochastic capacity expansion literature has been short in exploiting the benefits of parallel computing. In this work, we bridge this methodological gap by proposing a parallelization strategy for the considered algorithms, which allows us to arrive at solutions for the stochastic formulation of EVA within a few hours of computation.

C. Organization & contributions

(i) In terms of methodological contribution, we solve the stochastic EVA considering all available uncertainty conditions. (ii) Our algorithmic contribution amounts to proposing a parallel computing subgradient-based method and a novel parallel computing second-stage relaxation scheme, which decreases computational burden significantly. These algorithms are benchmarked against the commonly used Benders decomposition (which is also the algorithm that the PLEXOS commercial software uses). (iii) Our policy analysis contribution amounts to further testing the quality of the solution by comparing it against the ERAA 2021 deterministic approach.

The organization of the work is as follows. The EVA is formulated as a two-stage stochastic capacity expansion problem in section III. This formulation leads to a large-scale stochastic programming problem. A customized algorithm for tackling such a problem is proposed in section IV. Section V proposes parallel schemes for the described algorithms. Finally, the results are discussed in section VI.

II. EUROPEAN RESOURCE ADEQUACY ASSESSMENT

The European resource adequacy assessment proposes a methodology for measuring the ability of the power system

¹Note that transmission expansion is out of scope for the EVA, and thus we do not concern ourselves with this aspect in the present paper.

²PSR is a consulting firm based in Rio de Janeiro that has pioneered the commercialization of the SDDP algorithm for hydro-thermal planning

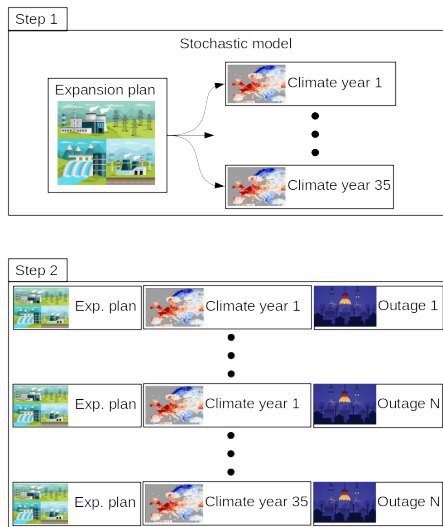


Fig. 1. The ERAA methodology consists of 2 steps. The first step, the so-called economic viability assessment, is calculated using a set of uncertain climatic conditions. The second step computes adequacy indicators, which use the previously computed expansion plan over an uncertainty set that consists of climatic conditions as well as random outages.

to react to uncertain events [44]. The overall methodology consists of two main blocks. The first block aims at determining investment and retirement opportunities (which we refer to as the expansion plan), the so-called Economic Viability Assessment (EVA). The second block uses these opportunities in order to measure adequacy indicators.

Figure 1 provides a visual representation of the overall ERAA methodology. During the first step (the EVA), a variety of uncertain climate conditions are considered (ENTSO-E considers a total of 35 climatic years for ERAA 2021 and ERAA 2022). The objective is to decide on investment and retirement that minimize the expected operational cost of the system. Note that this has to be decided before the realization of uncertainty, thus leading to a stochastic problem. During the second step, the adequacy indicators are calculated. The previously computed expansion plan is used for this purpose, and it is evaluated over an uncertainty set that consists of climate years as well as random outage patterns.

We highlight a critical difference between these steps. Step 2 can be decoupled into several independent problems, one for each climate year and outage pattern. This implies that it is computationally tractable. The situation regarding step 1 is considerably more involved, as it naturally links all climate years into a single problem, meaning that it is not possible to decouple the climatic years into independent problems. In our work, we focus on the first step. In order to tackle it, we propose a parallel computing algorithmic framework.

Due to increased computational complexity, ENTSO-E has considered simplified approaches in order to tackle step 1. For ERAA 2021 the simplification has been two-fold. (i) On the one hand, a scenario reduction methodology is applied which selects 7 out of 35 climatic years. (ii) The optimal expansion plan for each one of the selected climatic years is calculated, thus obtaining 7 expansion plans. The average between these



Fig. 2. The pan-European power system modelled in the ERAA 2021.

expansion plans is selected as the approximate solution of the stochastic problem. For ERAA 2022, a stochastic model is proposed. However, due to increased computational complexity, the model is reduced to 3 out of 35 climatic years, which are selected using a scenario reduction technique.

ENTSO-E has provided access to the ERAA 2021 input data for our team in the context of this work. For ERAA 2021, the EVA is modeled as a two-stage problem: the expansion plan is decided for a single target year. Note that ERAA 2022 adopts a multi-stage approach, namely there are consecutive years and the expansion is decided just before each year. In this work, we focus on ERAA 2021, which is the more natural step, since not even this model can be tackled in its stochastic form by state-of-the-art solvers.

The ERAA 2021 data considers the entire interconnected pan-European network, which is represented by considering 56 bidding zones that correspond to 36 countries. This is depicted in Figure 2. In terms of generators, the data contains the installed capacity mix per zone. We highlight that the generator data is aggregated per technology, per zone. A detailed overview of the installed capacity mix per zone can be found in [45]. The generator data also contains time series of planned maintenance of generators, de-rating factors, and must-run profiles of the generators. The data represents the transmission network as a transportation network, and includes transmission lines between zones, together with their limits. The data contains several candidate retirement opportunities, per zone, for thermal generators. Two thermal investment candidates, per zone, are considered. The uncertainty arises in the form of the so-called climatic years, each one being a time series (per zone) of demand, PV, wind, and inflow profiles. ENTSO-E has considered 35 climatic years for the EVA of the ERAA 2021 edition, which are depicted in Fig. 3.

ENTSO-E uses the ERAA 2021 data as input for their modeling tool for compiling the EVA expansion problem. As this tool is unable to tackle the problem in its stochastic form, we instead use the exact same input data in order to put together an open-source Julia [46] version, which enables us to develop a decomposition scheme for solving the problem, using the full set of 35 climatic years. We highlight that both models are built with the same modelling specifications.

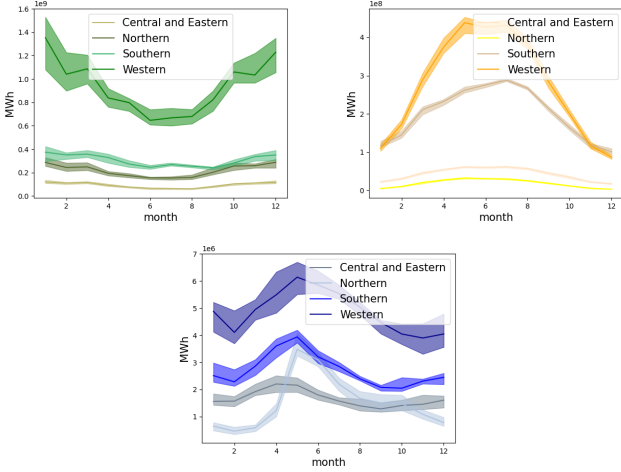


Fig. 3. An overview of climatic years of the 4 main European regions. The upper and lower bounds represent the 0.75 and 0.25 quantiles. The upper left panel presents wind production, the upper right is PV production, while the lower figure corresponds to hydrological inflows.

In particular, similarly to the EVA ENTSO-E model, ours is a linear model which takes into account the installed capacity mix per zone for proposing investment and retirement decisions (meaning that it is not a greenfield model).

III. EXPANSION PROBLEM

The stochastic capacity expansion problem is formulated as a two-stage stochastic program [18]. The first stage determines investments and retirements of technologies. The second stage solves an economic dispatch over a target year. The introduction of uncertainty takes place during the second stage. Each uncertainty realization corresponds to a so-called climatic year, which consists of a time series of demand, solar production, wind production, and hydro inflows. In order to ease the exposition, the sets, variables, and parameters of the problem are presented in the appendix. Furthermore, parameters are denoted in upper case while optimization variables are in lower case.

A. Expansion constraints

There is a maximum amount of plausible invested/retired capacity. This is modeled with the following constraints:

$$x_{n,g} \leq X_{n,g} \text{ for all } n \in \mathcal{N}, g \in \mathcal{G} \quad (1)$$

$$x_{n,g}^* \leq X_{n,g}^* \text{ for all } n \in \mathcal{N}, g \in \mathcal{G}^* \quad (2)$$

B. Generator constraints

Given $\omega \in \Omega$, the minimum and maximum power generation capabilities of units are described by the following constraints:

$$p_{n,g,t,\omega}^* \leq x_{n,g}^* \text{ for all } n \in \mathcal{N}, g \in \mathcal{G}^*, t \in \mathcal{T}, \omega \in \Omega \quad (3)$$

$$p_{n,g,t,\omega} \leq P_{t,n,g}^{max} - x_{n,g} \quad (4)$$

$$P_{t,n,g}^{min} \leq p_{n,g,t,\omega} \quad (5)$$

$$\text{for all } n \in \mathcal{N}, g \in \mathcal{G}, t \in \mathcal{T}$$

Constraints 3, 4 model the power production of new and existing capacity. Constraint 5 models must-run obligations.

C. Transmission network

The transmission network is modeled as a transportation network. The constraints are then bounds on the transfer capacity. Given $\omega \in \Omega$, these can be described by the following constraints:

$$L_{n,l}^{min} \leq f_{n,l,t,\omega} \leq L_{n,l}^{max} \quad (6)$$

for all $n \in \mathcal{N}, l \in \mathcal{L}(n), t \in \mathcal{T}$

Transportation models fail to accurately represent the true physics of power flow in European network models. This has motivated the introduction of a zonal PTDF approximation in European market clearing models. This implies the addition of linear constraints, which do not affect the overall structure of the decomposition schemes that are proposed in this paper.

D. Batteries

Batteries are modeled as energy storage resources available per zone. For a given $\omega \in \Omega$, they are modeled as follows:

$$bv_{n,t,\omega} = bv_{n,t-1,\omega} + BCE \cdot bc_{n,t,\omega} - BDE \cdot bd_{n,t,\omega} \quad (7)$$

$$bv_{n,t,\omega} \leq BV \quad (8)$$

$$bc_{n,t,\omega} \leq BC \quad (9)$$

$$bd_{n,t,\omega} \leq BD \quad (10)$$

$$\text{for all } n \in \mathcal{N}, t \in \mathcal{T}$$

Constraint 7 models the load balance of the battery. Constraints 8, 9, 10 model the maximum capacity, charge, and discharge capabilities of the battery.

E. Hydro power

Hydropower generation is modeled using four different hydro technologies: run-of-river, reservoir, pumped storage open loop, and pumped storage closed loop. The hydrology is simplified by ENTSO-E by considering the aggregated hydrological capabilities of each zone. The inflows are measured as equivalent energy inflows (MW). For a given $\omega \in \Omega$, each one of these technologies is modeled as follows. For ease of exposition, a variable associated with a certain technology has an associated subscript: run-of-river R , reservoir S , pumped storage open loop O , and pumped storage closed loop C .

- Run-of-River. There is no storage capability, therefore the net inflows are considered as turbined water.

$$q_{n,R,t,\omega} = A_{n,R,t,\omega} \text{ for all } n \in \mathcal{N}, t \in \mathcal{T} \quad (11)$$

- Reservoir: There is storage capability, consequently the water can be turbined or stored.

$$v_{n,S,t,\omega} = v_{n,S,t-1,\omega} + A_{n,S,t,\omega} - q_{n,S,t,\omega} - s_{n,S,t,\omega} \quad (12)$$

$$v_{n,S,t,\omega} \leq V_{n,S} \quad (13)$$

$$q_{n,S,t,\omega} \leq Q_{n,S} \quad (14)$$

$$\text{for all } n \in \mathcal{N}, t \in \mathcal{T}$$

Constraint 12 describes the reservoir dynamics, while constraints 13, 14 describe the maximum storage and maximum power production capacities respectively.

- Pumped storage open loop: There are head and tail reservoirs. The head turbines water to the tail reservoir, thus producing power. In low-demand periods the tail pumps water back to the head reservoir, thus consuming power. The system is exposed to natural inflows.

$$v_{n,O,t,\omega}^H = v_{n,O,t-1,\omega}^H + A_{n,O,t,\omega} + PE \cdot d_{n,O,t,\omega} - q_{n,O,t,\omega} - s_{n,O,t,\omega} \quad (15)$$

$$v_{n,O,t,\omega}^T = v_{n,O,t-1,\omega}^T - PE \cdot d_{n,O,t,\omega} + q_{n,O,t,\omega} \quad (16)$$

$$v_{n,O,t,\omega}^H \leq V_{n,O} \quad (17)$$

$$q_{n,O,t,\omega} \leq Q_{n,O} \quad (18)$$

$$d_{n,O,t,\omega} \leq D_{n,O} \quad (19)$$

$$\text{for all } n \in \mathcal{N}, t \in \mathcal{T}$$

Constraints 15, 16 are the water balance constraints of the head and tail reservoirs respectively. Note that the head reservoir is subject to rainfall uncertainty. Constraint 17 bounds the maximum storage, while constraints 18, 19 bound the turbined and pumped water respectively.

- Pumped storage closed loop: this is modeled in the same way as the open loop case, with the difference that no natural inflows are considered.

$$v_{n,C,t,\omega}^H = v_{n,C,t-1,\omega}^H + PE \cdot d_{n,C,t,\omega} - q_{n,C,t,\omega} - s_{n,C,t,\omega} \quad (20)$$

$$v_{n,C,t,\omega}^T = v_{n,C,t-1,\omega}^T - PE \cdot d_{n,C,t,\omega} + q_{n,C,t,\omega} \quad (21)$$

$$v_{n,C,t,\omega}^H \leq V_{n,C} \quad (22)$$

$$q_{n,C,t,\omega} \leq Q_{n,C} \quad (23)$$

$$d_{n,C,t,\omega} \leq D_{n,C} \quad (24)$$

$$\text{for all } n \in \mathcal{N}, t \in \mathcal{T}$$

F. Load balance

The load balance constraint aims at satisfying the demand of each zone using the resources of each zone plus imports from neighbouring zones. For a given $\omega \in \Omega$, it is formulated as follows:

$$\begin{aligned} p_{n,g,t,\omega} + p_{n,g,t,\omega}^* + bd_{n,t,\omega} + \sum_{h \in \mathcal{H}(n)} q_{n,h,t,\omega} + \sum_{l \in \mathcal{L}(n)} f_{n,l,t,\omega} \\ + ls_{n,t,\omega} + \mathcal{PV}_{n,t,\omega} + \mathcal{W}_{n,t,\omega} \\ = \mathcal{D}_{n,t,\omega} + ps_{n,t,\omega} + bc_{n,t,\omega} + \sum_{r \in \{C,O\}} d_{n,r,t,\omega} \end{aligned} \quad (25)$$

$$\text{for all } n \in \mathcal{N}, t \in \mathcal{T}$$

We highlight that reserve requirements are modelled in the ERAA 2021 study as extra load. More accurate models for reserve requirements have been studied in the past [47], their inclusion implies the addition of linear constraints which do not disrupt the overall structure of the algorithms developed in this paper.

G. Objective function

For a given $\omega \in \Omega$, we define the following quantities:

$$\sum_{n \in \mathcal{N}, g \in \mathcal{G}^*} IC_g \cdot x_{n,g}^* + FOM_g \cdot x_{n,g}^* - \sum_{n \in \mathcal{N}, g \in \mathcal{G}} FOM_g \cdot x_{n,g} \quad (26)$$

$$\sum_{n \in \mathcal{N}, g \in \mathcal{G}, t \in \mathcal{T}} VOM_g \cdot p_{n,g,t,\omega} + FC_g \cdot p_{n,g,t,\omega} \quad (27)$$

$$\sum_{n \in \mathcal{N}, g \in \mathcal{G}^*, t \in \mathcal{T}} VOM_g \cdot p_{n,g,t,\omega}^* + FC_g \cdot p_{n,g,t,\omega}^* \quad (28)$$

$$\sum_{n \in \mathcal{N}, l \in \mathcal{L}(n), t \in \mathcal{T}} f_{n,l,t,\omega} \cdot WC \quad (29)$$

$$\sum_{n \in \mathcal{N}, h \in \mathcal{H}(n), t \in \mathcal{T}} s_{n,h,t,\omega} \cdot SC \quad (30)$$

$$\sum_{n \in \mathcal{N}, t \in \mathcal{T}} ls_{n,t,\omega} \cdot VOLL \quad (31)$$

Eq. (26) is the first-stage cost. It consists of the investment cost plus the fixed maintenance cost of new capacity $x_{n,g}^*$, minus the fixed maintenance cost of retired capacity $x_{n,g}$. Eq. (27) corresponds to the cost of producing $p_{n,g,t,\omega}$ units of power, similarly Eq. (28) corresponds to the generation cost of new capacity. Eq. (29) corresponds to the cost of transporting power through the transmission network. Eq. (30) corresponds to a penalty for water spillage. Finally, Eq. (31) is the cost of involuntary load shedding, which is penalized at VOLL. Putting these elements together leads to the stochastic capacity expansion problem, which is described as follows:

$$\begin{aligned} \min_{x, x^*} (26) + \mathbb{E}_\omega \left[\min_{p_\omega, p_\omega^*, f_\omega, ls_\omega} (27) + (28) + (29) + (30) + (31) \right] \\ \text{s.t. (1) - (2)} \\ (3) - (25) \text{ for all } \omega \in \Omega \end{aligned}$$

(CEP)

Note that we have distinguished between two sets of constraints. The former refers to the so-called first-stage constraints, therefore they do not depend on $\omega \in \Omega$. The latter set of constraints is the second-stage constraints, they do depend on uncertainty, and there is one such set of constraints for each $\omega \in \Omega$. Note that this formulation implies that the first-stage variables $x_{n,g}, x_{n,g}^*$ are decided before uncertainty realizes, and thus do not depend on $\omega \in \Omega$.

IV. SOLUTION STRATEGY

The stochastic capacity generation expansion literature has often relied on Bender's decomposition (also known as the L-Shaped method when uncertainty is introduced) [18]. This scheme performs poorly as the number of expansion/retirement possibilities increases. In view of this, we propose a subgradient algorithm that is better suited for such situations. This technique allows us to reduce the number of iterations. However, it can still be computationally costly as the calculation of each subgradient is non-trivial. Consequently, we further propose a relaxation of the economic dispatch, which allows us to calculate subgradient approximations efficiently. The approximation is refined throughout iterations,




Fig. 4. Decomposition of the two-stage stochastic problem. Each node represents an uncertainty realization and has associated with it an optimization problem that aims at minimizing the costs of the stage.

thus ensuring convergence. Such a technique allows us to reduce the computational burden significantly. We begin by briefly describing the L-shaped technique, followed by our algorithmic contributions.

The L-shaped method breaks down the overall problem into smaller subproblems, first by considering separate subproblems for the first and second stages, and secondly by considering a subproblem for each uncertainty realization for the second stage, as presented in Figure 4. Given an uncertainty realization ω and a first-stage decision x, x^* , the second-stage subproblem for a given ω is as follows:

$$\begin{aligned} \mathcal{V}(x, x^*, \omega) = & \min_{p_\omega, P_\omega^*, f_\omega, l_{S\omega}} (27) + (28) + (29) + (30) + (31) \\ & \text{s.t. (3)} \quad (\lambda_{n,g,t,\omega}^*) \\ & \quad \quad (4) \quad (\lambda_{n,g,t,\omega}) \\ & \quad \quad (5) - (25) \end{aligned} \quad (31)$$

where $\lambda_{n,g,t,\omega}^*, \lambda_{n,g,t,\omega}$ are dual multipliers of constraints (3), (4) respectively. These subproblems allow us to re-write the CEP problem as

$$\begin{aligned} & \min_{x, x^*} (26) + \mathbb{E}_\omega [\mathcal{V}(x, x^*, \omega)] \\ & \text{s.t. (1) - (2)} \end{aligned} \quad (\text{CEP-R})$$

The function $\mathbb{E}_\omega [\mathcal{V}(x, x^*, \omega)]$ is piece-wise linear convex in x, x^* [18], and can therefore be under-approximated by supporting hyperplanes, commonly referred as to cuts. These cuts are computed by calculating the subgradients of the second-stage functions $\mathcal{V}(x, x^*, \omega)$ [18]. Consequently, given a collection of supporting hyperplanes $\{c_i(x, x^*)\}_{i=1}^N$, the previous problem can be approximated as follows:

$$\begin{aligned} & \min_{x, x^*, \theta} (26) + \theta \\ & \text{s.t. (1) - (2)} \\ & \quad \theta \geq c_i(x, x^*) \text{ for all } i \in 1, \dots, N \end{aligned} \quad (\text{M})$$

The L-Shaped algorithm advances by finding, at each iteration, a new supporting hyperplane for problem M. Each iteration begins by solving problem M which leads to a trial action \hat{x}^i, \hat{x}^{*i} . The second-stage subproblems $\mathcal{V}(x, x^*, \omega)$ are then solved around \hat{x}^i, \hat{x}^{*i} . The dual multipliers $\lambda_{n,g,t,\omega}, \lambda_{n,g,t,\omega}^*$ are a subgradient of $\mathcal{V}(x, x^*, \omega)$ at \hat{x}^i, \hat{x}^{*i} , and can therefore be used for computing a supporting hyperplane [18],

which is added to problem M. As we move through iterations, the method starts building an accurate representation of $\mathbb{E}_\omega [\mathcal{V}(x, x^*, \omega)]$ around the optimal region, eventually finding the optimal value. In fact, the method converges after finitely many iterations [48].

A. Subgradient algorithm

A drawback of the L-Shaped scheme is that, as the dimensionality of x, x^* increases, additional supporting hyperplanes are required in order to describe $\mathbb{E}_\omega [\mathcal{V}(x, x^*, \omega)]$. As a consequence, finding the optimal region may require many extra costly iterations. Instead, in this work we use a subgradient algorithm. Given an initial expansion candidate \hat{x}, \hat{x}^* , we specifically calculate a subgradient of the objective function of the CEP-R problem around \hat{x}, \hat{x}^* and update the candidate expansion plan along the direction of such a subgradient. This method has the following advantages: (i) it does not require a hyperplane description of $\mathbb{E}_\omega [\mathcal{V}(x, x^*, \omega)]$ in order to advance to the next candidate \hat{x}, \hat{x}^* ; (ii) it can be initialized around a trial \hat{x}, \hat{x}^* which is known in advance to be somewhat close to the optimal value, thus ensuring that if the starting value is close to the optimal solution the iterates remain in the optimal region and will need few iterations.

We start by decomposing problem CEP by rewriting it as CEP-R. Note that a subgradient of the objective function along the $x_{n,g}^*$ coordinate is: $\rho_{n,g}^* = IC_{n,g} + FOM_{n,g} + \mathbb{E}_\omega \left[\sum_{t \in \mathcal{T}} \lambda_{n,g,t,\omega}^* \right]$. And a subgradient along the $x_{n,g}$ coordinate is: $\rho_{n,g} = -FOM_{n,g} + \mathbb{E}_\omega \left[\sum_{t \in \mathcal{T}} \lambda_{n,g,t,\omega} \right]$.

Due to the reformulation CEP-R, the calculation of these slopes can be decomposed into calculating several subproblems, concretely by solving $\mathcal{V}(x, x^*, \omega)$ for each $\omega \in \Omega$. This decomposition allows us to apply the following scheme. We commence by providing an initial candidate action \hat{x}^1, \hat{x}^{*1} . During each iteration, the subproblems $\mathcal{V}(\hat{x}^i, \hat{x}^{*i}, \omega)$ are solved for all $\omega \in \Omega$. Using the dual multipliers of these subproblems, the subgradients $\rho_{n,g}, \rho_{n,g}^*$ are calculated. Finally, the trial action is updated through a projected subgradient step: $\hat{x}_{n,g}^{i+1} = \min\{X_{n,g}, \hat{x}_{n,g}^i + \alpha^i \cdot \rho_{n,g}^*\}$, $\hat{x}_{n,g}^{*i+1} = \min\{X_{n,g}^*, \hat{x}_{n,g}^{*i} + \alpha^i \cdot \rho_{n,g}\}$. The term α^i is a stepsize that is crucial to the performance of the algorithm [49]. We have selected the Polyak stepsize [49], which ensures convergence. The Polyak stepsize can be described as follows: $\alpha^i = \frac{W^* - W^i}{\sum (p_{n,g}^i)^2 + \sum (p_{n,g}^{*i})^2}$. Here, W^* is the optimal value of the CEP problem, while W^i is the current iterate objective value. As the optimal value W^* is not known in advance, an approximate value can be used. No lower bound is obtained, thus its stabilization is used as a stopping criterion.

B. Second-stage relaxation algorithm

Each iteration of the subgradient scheme can be costly because it carries the computational burden of solving each second-stage subproblem. For this reason, we propose a scheme that relaxes the second stage, namely the economic dispatch. Using such a relaxation, each iterate of a trial \hat{x}^i, \hat{x}^{*i} becomes efficient, thus allowing us to increase the search




Fig. 5. Chronological decomposition. The second stage is partitioned into several consecutive chronological blocks.

speed. The relaxation can then be refined in order to tighten the search and ensure convergence.

Let us begin by describing the second-stage relaxation. To achieve this, we resort to dynamic programming [50], and partition the second-stage horizon $1, \dots, \mathcal{T}$ into K consecutive blocks: $\{1, t_1\}, \{t_1 + 1, t_2\}, \dots, \{t_{K-1} + 1, \mathcal{T}\}$. This leads to the representation shown in Figure 5, where the second stage has been partitioned into several blocks. The subproblems at each block are given by the dynamic programming equations. The equation at block k is given by:

$$\begin{aligned} \mathcal{V}_k(x, x^*, bv_{t_{k-1}, \omega}, v_{t_{k-1}, \omega}, \omega) = \\ \min \left[(27) + (28) + (29) + (30) + (31) \right]_{t=t_{k-1}+1}^{t_k} \\ + \mathcal{V}_{k+1}(x, x^*, bv_{t_k, \omega}, v_{t_k, \omega}, \omega) \\ \text{s.t. } \left[(3) - (25) \right]_{t=t_{k-1}+1}^{t_k} \end{aligned}$$

Here, the notation indicates that the objective function and the constraints are restricted to $t = t_{k-1} + 1, \dots, t_k$. Note that, at the initial time boundary, i.e. at $t = t_{k-1} + 1$, constraint (7) requires the battery state of charge at $t = t_{k-1}$. Consequently, the notation indicates that \mathcal{V}_k depends on $bv_{t_{k-1}, \omega}$, the battery state of charge at stage t_{k-1} . Similarly, Eqs. (12), (15), (16), (20), (21) require the reservoir level at stage t_{k-1} , and this information is summarized in the vector $v_{t_{k-1}, \omega}$. At the final time boundary at $t = t_k$, the objective function includes \mathcal{V}_{k+1} , which captures the future costs of the system. Note that the subproblem of the first block satisfies $\mathcal{V}_1(x, x^*, \omega) = \mathcal{V}(x, x^*, \omega)$.

Each function \mathcal{V}_{k+1} is piecewise linear convex in $x, x^*, bv_{t_k, \omega}, v_{t_k, \omega}$ [18], and so we can approximate it using supporting hyperplanes. Thus, given a collection of supporting hyperplanes $\{c_i(x, x^*, bv_{t_k, \omega}, v_{t_k, \omega})\}_{i=1}^N$ an approximation of the subproblem at the k -th block is given by:

$$\begin{aligned} \widehat{\mathcal{V}}_k(x, x^*, bv_{t_{k-1}, \omega}, v_{t_{k-1}, \omega}, \omega) = \\ \min \left[(27) + (28) + (29) + (30) + (31) \right]_{t=t_{k-1}+1}^{t_k} + \theta_{k+1, \omega} \\ \text{s.t. } \left[(3) - (25) \right]_{t=t_{k-1}+1}^{t_k} \\ \theta_{k+1, \omega} \geq c_i(x, x^*, bv_{t_k, \omega}, v_{t_k, \omega}), i = 1 \dots N \end{aligned}$$

In particular, the approximation of the first-block subproblem $\widehat{\mathcal{V}}_1$ is an approximation of the second-stage subproblem,




Fig. 6. Step 2) of the second-stage relaxation algorithm. Step 2.1) solves the current approximation of problem CEP. Step 2.2) refines the approximation of problem CEP around the trial expansion plan found during step 2.1).

namely the economic dispatch. Consequently, we can approximate problem CEP-R as follows:

$$\begin{aligned} \min_{x, x^*} (26) + \mathbb{E}_\omega \left[\widehat{\mathcal{V}}_1(x, x^*, \omega) \right] \\ \text{s.t. } (1) - (2) \end{aligned} \quad (\text{CEP-A})$$

Calculating each $\widehat{\mathcal{V}}_1$ is straightforward, therefore problem CEP-A can be solved efficiently. Note, however, that, due to the approximation, the subgradients we may obtain are not necessarily tight. Consequently, the approximation is tightened throughout iterations, thus ensuring convergence. An initial approximation of $\mathcal{V}_1(x, x^*, \omega)$ is calculated at the beginning of the algorithm. The objective of doing so is to, similarly to the subgradient scheme, have an initial candidate expansion plan to start the search. These ideas are the basis of our algorithmic scheme, which is depicted in pseudo-code in algorithm 1.

The second-stage relaxation algorithm is illustrated graphically in Fig. 6. The figure presents step 2) of the algorithm. For ease of exposition, step 1), where a warm-start is calculated, is described afterward. Step 2) is subdivided into two steps. Step 2.1) focuses on the first node and the first-time step/nodes of the second stage. This corresponds to problem CEP-A. The algorithm uses the current approximation of $\mathcal{V}_1(x, x^*, \omega)$ to solve problem CEP-A (which approximates CEP) and obtain a candidate expansion plan $\widehat{x}^i, \widehat{x}^{*i}$. The algorithm proceeds with step 2.2), which aims at refining the approximation of $\mathcal{V}_1(x, x^*, \omega)$ around $\widehat{x}^i, \widehat{x}^{*i}$. To do so, the algorithm performs a forward pass (step 2.2.1) and a backward pass (step 2.2.2). The forward pass proceeds forward in the number of second-stage nodes, solving the first-node subproblem and proceeding until arriving at the last node. The algorithm continues with the backward pass. This step computes supporting hyperplanes around the storages found during the forward pass and around the trial expansion plan $\widehat{x}^i, \widehat{x}^{*i}$. Starting from the last second-stage node, the subproblem is solved, using the dual multipliers for estimating a supporting hyperplane for the subproblem of the preceding node. The process is repeated until reaching the node associated with the first node of the second stage. The algorithm performs this procedure for all uncertainty realizations. Having described step 2), we can now describe the warm-start of step 1). Given an initial candidate expansion plan $\widehat{x}^0, \widehat{x}^{*0}$ the objective is to provide an approximation of $\mathcal{V}_1(x, x^*, \omega)$ around the given initial point. To do so, the warm-start performs step 2.2) throughout several iterations. These several iterations are performed in order to ensure that the

storages found during the forward pass provide a reasonable approximation.

Remark IV.1. Note that both steps 2.1) and 2.2) can be solved efficiently. The former is a relatively small problem, which can be solved either by Bender's decomposition or subgradient schemes. In this work, we use Bender's decomposition. The latter step involves solving several small subproblems and thus does not pose a computational burden. As a consequence, this algorithmic approach is able to perform far more iterations than the subgradient algorithm.

Upper and lower bounds can be obtained as follows. An upper bound is computed by using the optimal values found during step 2.1) and during step 2.2.1). A lower bound can be obtained by the solution of step 2.1). By comparing upper and lower bounds, one can measure the optimality gap which can be used as a stopping criterion. The convergence of the decomposition algorithm is guaranteed due to the following lemma.

Algorithm 1 Second-stage relaxation algorithm

INPUT: Provide a lower bound for $\theta_{k,\omega}$ for $k = 1, \dots, K$, $\omega \in \Omega$, and an initial trial action \hat{x}^0, \hat{x}^{*0} .

- 1) WARM-START: Calculate an initial approximation of \mathcal{V}_1 for all $\omega \in \Omega$, around \hat{x}^0, \hat{x}^{*0} .
 - 2) **for** $i = 1, \dots, N$.
 - (2.1) Solve CEP-A using the current approximation of \mathcal{V}_1 . Thus, obtaining trial action \hat{x}^i, \hat{x}^{*i} .
 - (2.2) **for** $\omega \in \Omega$:
 - (2.2.1) **Forward Pass.** **for** $k = 1, \dots, K$:
Solve the approximated problem $\hat{\mathcal{V}}_k(\hat{x}^i, \hat{x}^{*i}, \hat{bv}_{t_{k-1},\omega}^i, \hat{v}_{t_{k-1},\omega}^i, \omega)$ and get the optimal storage $\hat{bv}_{t_k,\omega}^i, \hat{v}_{t_k,\omega}^i$.
 - (2.2.2) **Backward Pass.** **for** $k = K, \dots, 2$:
 - (2.2.2.1) Solve the approximation problem $\hat{\mathcal{V}}_k(\hat{x}^i, \hat{x}^{*i}, \hat{bv}_{t_{k-1},\omega}^i, \hat{v}_{t_{k-1},\omega}^i, \omega)$.
 - (2.2.2.2) Using the dual multipliers compute a supporting hyperplane around $\hat{x}^i, \hat{x}^{*i}, \hat{bv}_{t_{k-1},\omega}^i, \hat{v}_{t_{k-1},\omega}^i$.
 - (2.2.2.3) Add the supporting hyperplane to problem $\hat{\mathcal{V}}_{k-1}(x, x^*, bv_{t_{k-2},\omega}, v_{t_{k-2},\omega}, \omega)$.
-

Lemma IV.2. Algorithm 1 converges in a finite number of iterations to CEP.

Proof. We prove that if no new cuts are added to $\hat{\mathcal{V}}_1(x, x^*, \omega)$ for $\omega \in \Omega$, then we are at the optimal of CEP. Suppose that no new cuts are obtained after iteration i and consider the expansion plan at that iteration x^i, x^{*i} . Note that, if no new cuts are added, then $\hat{\mathcal{V}}_1(x^i, x^{*i}, \omega) = \mathcal{V}_1(x^i, x^{*i}, \omega)$ (otherwise during the $i + 1$ iteration we would have found a new cut). Now let us assume that we have not converged, therefore there exists \bar{x}, \bar{x}^* for which $(26) + \mathbb{E}_\omega[\mathcal{V}_1(\bar{x}, \bar{x}^*, \omega)] < (26) + \mathbb{E}_\omega[\mathcal{V}_1(x^i, x^{*i}, \omega)]$. CEP-A is an under approximation of CEP.

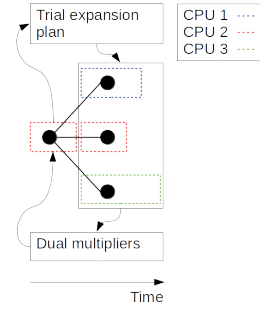


Fig. 7. Parallel subgradient algorithm. The master CPU provides a trial expansion plan, which is sent to the subproblems. The subproblems are distributed among the CPUs and solved, the dual multipliers are sent to the master CPU.

Thus, its maximum possible value is the left-hand side of the previous inequality. As a consequence,

$$\begin{aligned}
 (26) + \mathbb{E}_\omega[\hat{\mathcal{V}}_1(x^i, x^{*i}, \omega)] &< (26) + \mathbb{E}_\omega[\mathcal{V}_1(\bar{x}, \bar{x}^*, \omega)] \\
 &< (26) + \mathbb{E}_\omega[\mathcal{V}_1(x^i, x^{*i}, \omega)] \\
 &= (26) + \mathbb{E}_\omega[\hat{\mathcal{V}}_1(x^i, x^{*i}, \omega)]
 \end{aligned}$$

This leads to a contradiction, therefore we conclude that x^i, x^{*i} is optimal.

Finally, let us show that this can be achieved in finite iterations. Following the same idea as in [48], where it is shown that SDDP-type [36] algorithms converge, we start from the subproblem of the last block. Note that the set of dual multipliers of the subproblem associated to $\mathcal{V}_K(x, x^*, bv_{t_{K-1},\omega}, v_{t_{K-1},\omega}, \omega)$ corresponds to the vertices of the feasibility set of the dual problem, and it does not depend on $x, x^*, bv_{t_{K-1},\omega}, v_{t_{K-1},\omega}$. Consequently, there are finitely many supporting hyperplanes for the last subproblem, and thus after finitely many iterations we arrive at the supporting hyperplane description of $\mathcal{V}_K(x, x^*, bv_{t_{K-1},\omega}, v_{t_{K-1},\omega}, \omega)$. We can now use such a description for the subproblem at block $K - 1$, and apply the same argument. Inductively, we can proceed backward in the number of blocks. \square

V. PARALLEL SCHEME

The present section aims at describing a parallel scheme for the algorithms presented in section IV. In this work, we have considered synchronous parallel implementations.

• **Parallel subgradient algorithm:** The parallelization strategy followed for the subgradient algorithm is presented graphically in Fig. 7. Each iteration proceeds as follows. The master CPU provides an initial candidate expansion plan. The second-stage subproblems are distributed among the available CPUs. Each CPU solves its associated subproblem, the dual multipliers are collected and sent to the master CPU. At this point, the CPUs synchronize, i.e., a CPU stays idle until all other CPUs have finished their job. The master CPU uses the dual information in order to apply a projected subgradient and update the trial expansion plan.

- **Parallel L-shaped algorithm:** The parallelization follows a similar strategy as in the previous scheme. The difference is that during each iteration the master CPU calculates problem M .
- **Parallel second-stage relaxation algorithm:** The parallel scheme for our subgradient-relaxation algorithm follows a similar procedure. Step 2.1) is parallelized using the strategy described in Fig. 7. Step 2.2) is parallelized as follows. The uncertainty realizations are distributed among the available CPUs. Each CPU performs steps 2.2.1 and 2.2.2 of algorithm 1 (see Fig. 6). At this point, the CPUs synchronize, so that the fastest CPU waits for the slowest CPU.

VI. CASE STUDY

The EVA problem aims at determining the expansion/retirement plans that will occur, and our study considers the 2025 target year. We have considered a total of 12 blocks per day.

Our algorithms are implemented in Julia v1.5 and JuMP v22. The chosen linear programming solver is Gurobi 9. The computational work is performed on the Lemaitre3 cluster of UCLouvain, which is hosted at the Consortium des Equipements de Calcul Intensif (CECI). The cluster consists of 80 compute nodes with two 12-core Intel SkyLake 5118 processors at 2.3 GHz and 95 GB of RAM (3970MB/core), interconnected with an OmniPath network (OPA-56Gbps).

A. Value of the stochastic solution

Due to the scale of the model, ENTSO-E has considered an approximate solution in ERAA 2021. This approximation proceeds by solving the so-called wait-and-see solution [18] of problem CEP. This leads to a candidate expansion plan x_ω, x_ω^* , for each $\omega \in \Omega$. The average expansion plan \bar{x}, \bar{x}^* is used as an approximation of the stochastic expansion plan. A natural question is whether this is a good approximation. To measure this, we can use well-established bounds. The following relationship holds [18]:

$$\text{Wait-and-see} \leq \text{CEP} \leq (26) + \mathbf{E}_\omega \left[\mathcal{V}(\bar{x}, \bar{x}^*, \omega) \right]$$

Note that the right-hand side is the objective of CEP when using the sub-optimal average expansion plan \bar{x}, \bar{x}^* . We refer to this as the average-W.S. solution. Note that these bounds do not require the calculation of the stochastic solution, and thus provide a reasonable way to measure if a stochastic solution is of interest for the problem. The wait-and-see solution has a cost of $5.0220e10$ €, while the average wait-and-see solution has a cost of $5.2298e10$ €. As one can observe, the relative difference between both quantities is approximately 4.13%, and thus the stochastic solution stands to improve the deterministic approximation by at most this quantity. This difference is of interest: adequacy studies such as ERAA aim at capturing adequacy metrics that involve curtailments. These curtailments, calculated as in Eq. (31), represent less than 2% of the total costs, and thus a solution whose total costs differ in 4% is clearly of interest.




Fig. 8. The optimality gap evolution of the L-shaped method and our subgradient-relaxation algorithm. The x-axis is the elapsed time, while the y-axis is the relative difference between the upper and lower bounds.

B. Stochastic solution

The previous subsection presented the relevance of a stochastic solution. The present subsection discusses how to obtain such a solution using the previously discussed algorithms. The algorithms have been run using 35 CPUs. In the case of the subgradient relaxation, we have decomposed the second stage into 92 blocks.

We begin by examining the convergence evolution of the L-shaped method. To do so, we examine the optimality gap evolution, which is presented in Fig. 8. As one can observe, the L-shaped method struggles to close the optimality gap. In fact, after more than a day of computation, the obtained gap is not of practical use. On the other hand, our subgradient-relaxation scheme is able to provide an optimality gap near 1% after about 4 hours of computation.

The proposed subgradient algorithm does not provide a lower bound estimate, therefore the calculation of an optimality gap is not presented. Figure 9 presents the upper bound evolution of the subgradient algorithm and the subgradient-relaxation scheme. Recall that the subgradient-relaxation scheme uses an approximation of the economic dispatch, thus it does not correspond to the true value of using the expansion plan. Consequently, in order to have comparable upper bound values, after running the subgradient-relaxation algorithm we have evaluated the true cost of using the obtained expansion plan, this corresponds to the horizontal non-dashed green line. This quantity and the upper bound of the subgradient algorithm are comparable. The left panel of Fig. 9 presents the results when using 12 blocks per day, while the right panel presents the results when using 24 blocks per day. We note that both algorithmic schemes are able to converge. However, it can be observed that the subgradient-relaxation scheme converges considerably faster. In fact, for the 12-block, system after 30 hours of computation, the subgradient scheme has not attained the bound that the subgradient-relaxation scheme is able to find in just 4 hours. For the 24-block system, after almost 2 days of computation, the subgradient scheme fails to attain the bound that the subgradient-relaxation scheme finds in 15 hours. This indicates that the subgradient-relaxation scheme is better suited as the size of the problem increases. We note that the obtained run times can be used in practice. ENTSO-E's experience with a stochastic model, which consists of climate years and is solved as a large LP for a variety of target years, is that it requires longer run times.




Fig. 9. The convergence evolution of the subgradient algorithm and the subgradient-relaxation scheme. The x-axis is the elapsed time while the y-axis is the cost. The yellow horizontal lines correspond to the wait-and-see solution and ENTSO-E’s solution. The left panel corresponds to a setting with 12 blocks per day, while the right panel uses 24 blocks per day.




Fig. 10. Relative total cost comparison between average wait-and-see (Av. W.S.), Wait-and-See (W.S.), and Stochastic (St.) solutions.

Alternative approaches are considered for solving the problem, as a means of further comparison. This includes the progressive hedging algorithm [24] and Benders Decomposition With Multiple Master Problems (BDMM) [25]. Unfortunately, these schemes are not suitable for our setting. On the one hand, each subproblem of progressive hedging is harder than Benders as (i) it includes a quadratic term in the objective, and (ii) includes the optimization of the first and second stage. This renders each iteration prohibitively expensive. On the other hand, BDMM increases the amount of second-stage Benders subproblems. In the case of the problem that we are interested in this is not tractable, as each Benders subproblem is large.

We highlight that the use of parallel computing has been crucial in order to obtain the solutions that are indicated here. In fact, as each CPU is handling one of the 35 climatic years, we expect a serial run to be approximately 35 times more time-consuming.

C. Solution analysis

The present subsection aims at studying the differences between the stochastic solution obtained in subsection B and the deterministic solution obtained in subsection A. In terms of total costs, as presented in Fig. 10, we observe a difference of nearly 2.7% between the stochastic solution and the wait-and-see solution, a difference that supports the value of having computed a stochastic solution. Furthermore, we observe that the stochastic solution provides an improvement of nearly 1.5% with respect to the average W.S. solution, implemented by ENTSO-E in ERAA 2021.

This difference in total costs is examined in Fig. 11 by decomposing the total costs into operational costs, retired capacity savings, expansion plan costs, and curtailment costs. The upper panel presents the costs for both approaches, the stochastic solution, and the average W.S. solution, while the lower panel presents the relative difference between these

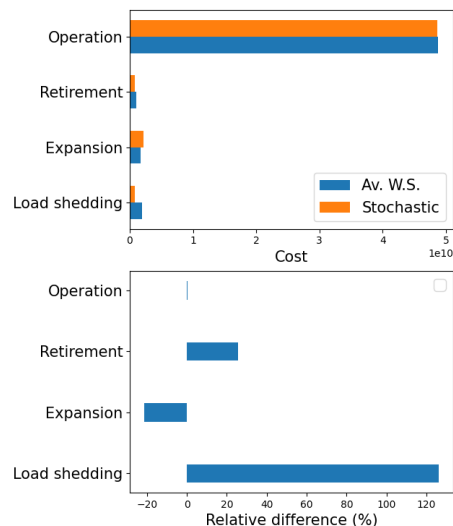


Fig. 11. Cost breakdown. The upper panel presents the total costs, while the lower panel presents the relative difference between the quantities shown in the upper panel.

solutions (a positive number indicates that the average W.S. solution has a higher value). Note that the stochastic solution tends to result in significantly fewer curtailments, which can be explained by the increased investments and fewer retirements, that is to say, we arrive at a more conservative solution. On the other hand, the lack of information regarding the future possible outcomes in the average expansion plan approach makes the solution prone to over-retiring and underestimating the required investments, which in turn implies that the power system is unable to satisfy the demand more effectively.

D. Adequacy metrics

One objective of the ERAA study is to measure the ability of the system to maintain secure levels of supply. Two adequacy metrics are of particular interest to ENTSO-E. The first one is the loss of load expectation (LOLE), defined as $LOLE = \mathbf{E}_\omega [LOL_\omega]$. Here, LOL_ω is the number of hours during which demand is not served for the climatic year $\omega \in \Omega$. The second metric is the expected value of energy not served (EENS), which is defined as $EENS = \mathbf{E}_\omega [ENS_\omega]$. Here, ENS_ω is the number of curtailments obtained in climatic year $\omega \in \Omega$. We consider the 35 climatic years that are used for formulating the stochastic model, thereby aiming to compare its performance against the approximate solution that is obtained by averaging the expansion plans (which is the approach used by ENTSO-E in ERAA 2021). Therefore, there is no out-of-sample testing. We calculate the metrics of interest and present the results in Fig. 12. We consider both the LOLE and the EENS for the 4 main European regions: Central and Eastern Europe, Western Europe, Southern Europe, and Northern Europe. Fig. 12 effectively demonstrates how the use of a stochastic solution is able to provide consistent and significantly more accurate metrics, which can lead to a better-informed adequacy assessment in Europe. This attribute is of particular interest to studies such as ERAA.

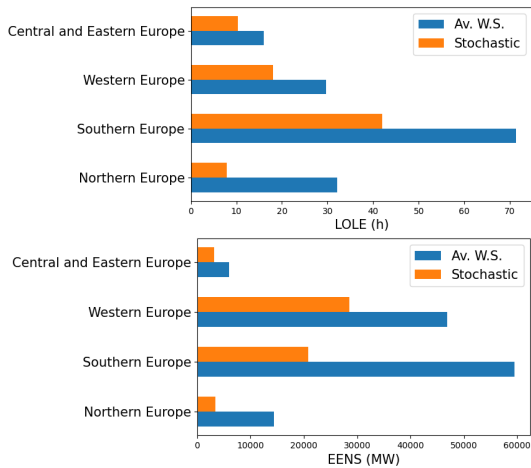


Fig. 12. The LOLE and EENS metrics for different European regions.

VII. CONCLUSIONS

In this work, we present a high-performance computing approach for obtaining a stochastic solution for the EVA of ERAA 2021. We specifically propose a subgradient algorithm implemented in parallel computing infrastructure, as well as a subgradient-relaxation approach that emerges as being practically attractive due to its capability to tackle large-scale problems efficiently. We further observe that the commonly used L-Shaped method is unable to provide a solution of practical use in a reasonable amount of time. Finally, we compare the stochastic solution against a deterministic approach implemented by ENTSO-E for the ERAA 2021 edition. A noticeable difference between both solutions in terms of commonly used adequacy metrics emerges, which highlights the practical value of the stochastic solution that we are able to compute.

Future ERAA studies aim at improving the EVA on two fronts. (i) On the one hand, a multi-year stochastic model is considered, that decides investments and retirements before the realization of each year. (ii) On the other hand, a more robust model is considered, which includes further sources of uncertainty, such as random outage patterns. These increased uncertainty motivates the study of scenario reduction techniques.

Furthermore, fuel price assumptions can have a significant impact on the results, and are also quite volatile, as the recent energy crisis suggests. The proposed approach can handle objective function uncertainty, thus allowing us to include these assumptions directly into the model. On the other hand, ENTSO-E's interaction with industry stakeholders points that this uncertainty can be considered as a sensitivity run of the model with different assumptions. This can provide insights on the impact without introducing further computational complexity to the model.

ACKNOWLEDGMENTS

Computational resources have been provided by the Consortium des Équipements de Calcul Intensif (CÉCI), funded by

the Fonds de la Recherche Scientifique de Belgique (F.R.S.-FNRS) under Grant No. 2.5020.11 and by the Walloon Region. This work was supported by computational time granted from the National Infrastructures for Research and Technology S.A. (GRNET S.A.) in the National HPC facility - ARIS - under project ID 013010. This project has received funding from the European Research Council (ERC) under the European Union Horizon 2020 research and innovation program (grant agreement No. 850540). This work was supported by the Research Fund of the Facultad de Ciencias, Universidad de los Andes INV-2021-128-2307. The research presented in this paper was made during the lead author's internship at ENTSO-E, for which the authors are grateful.

NOMENCLATURE

- **Sets:** Ω : The set of uncertainty realizations.
- \mathcal{N} : The set of zones.
- \mathcal{G} : The set of existing generators.
- \mathcal{G}^* : The set of new generators.
- \mathcal{T} : The time horizon of the economic dispatch.
- $\mathcal{L}(n)$: The set of lines for zone n .
- $\mathcal{H}(n)$: The set of hydro technologies for zone n .
- **Parameters:** IC: The annualized investment cost (€/MW).
- FOM: The fixed operational and maintenance cost (€/MW).
- VOM: The variable cost (€/MW).
- FC: The fuel costs (€/MW).
- WC: The wheeling charge cost (€/MW).
- SC: The spillage cost (€/MW).
- VOLL: A high cost for curtailment (€/MW).
- $X_{n,g}$: Retirement limit for $n \in \mathcal{N}, g \in \mathcal{G}$ (MW).
- $X_{n,g}$: Expansion limit for $n \in \mathcal{N}, g \in \mathcal{G}^*$ (MW).
- $P_{t,n,g}^{max}$: Maximum production, $t \in \mathcal{T}, n \in \mathcal{N}, g \in \mathcal{G}^*$ (MW).
- $P_{t,n,g}^{min}$: Minimum production, $t \in \mathcal{T}, n \in \mathcal{N}, g \in \mathcal{G}^*$ (MW).
- $X_{n,g}$: Expansion limit for $n \in \mathcal{N}, g \in \mathcal{G}^*$ (MW).
- $L_{n,l}^{max}$: Transfer limit, $n \in \mathcal{N}, l \in \mathcal{L}(n)$ (MW).
- $L_{n,l}^{min}$: Transfer minimum, $n \in \mathcal{N}, l \in \mathcal{L}(n)$ (MW).
- BCE: Battery charge efficiency (%).
- BDE: Battery discharge efficiency (%).
- BV: Battery capacity (MW).
- BC: Maximum charge capacity (MW).
- BD: Minimum discharge capacity (MW).
- $V_{n,h}$: Maximum storage for zone n , technology h (MWh).
- $Q_{n,h}$: Turbine capacity for zone n , technology h (MW).
- $D_{n,h}$: Pump capacity for zone n , technology h (MW).
- PE: Pump efficiency (%).
- **Parameters with uncertainty:** $A_{n,h,t,\omega}$: Inflow for $n \in \mathcal{N}, h \in \mathcal{H}(n), t \in \mathcal{T}, \omega \in \Omega$ (MW).
- $\mathcal{PV}_{n,t,\omega}$: PV production for $n \in \mathcal{N}, t \in \mathcal{T}, \omega \in \Omega$ (MW).
- $\mathcal{W}_{n,t,\omega}$: Wind production for $n \in \mathcal{N}, t \in \mathcal{T}, \omega \in \Omega$ (MW).
- $\mathcal{D}_{n,t,\omega}$: Demand for $n \in \mathcal{N}, t \in \mathcal{T}, \omega \in \Omega$ (MW).
- **Variables:** $x_{n,g}^*$: The invested capacity (MW).
- $x_{n,g}$: The retired capacity (MW).
- $p_{n,g,t,\omega} / p_{n,g,t,\omega}^*$: Power produced with existing/new generators (MW).
- $f_{n,l,t,\omega}$: Power transferred through lines (MW).
- $s_{n,h,t,\omega}$: The amount of spillage (MW).

$l_{s_{n,t,\omega}}$: The load shedding (MW).
 $bv_{n,t,\omega}$: The battery state of charge (MWh).
 $bc_{n,t,\omega}$: Power charged into the battery (MW).
 $bc_{n,t,\omega}$: The amount of charged power into the battery (MW).
 $bd_{n,t,\omega}$: The battery discharged power (MW).
 $q_{n,h,t,\omega}$: The hydro turbined produced power (MW).
 $v_{n,h,t,\omega}$: The state of storage of the reservoir (MWh).
 $d_{n,h,t,\omega}$: The pumped power (MW).

REFERENCES

- [1] "Clean energy package." https://energy.ec.europa.eu/topics/energy-strategy/clean-energy-all-europeans-package_en, accessed: 2022-09-26.
- [2] E. Commission, "Regulation (eu) 2019/941 of the european parliament and of the council of 5 june 2019 on risk-preparedness in the electricity sector and repealing directive 2005/89/ec," *Off. J. Eur. Union*, vol. 158, pp. 1–21, 2019.
- [3] —, "Regulation (eu) 2019/943 of the european parliament and of the council of 5 june 2019 on the internal market for electricity," *Off. J. Eur. Union*, vol. 158, pp. 54–124, 2019.
- [4] "European resource adequacy assessment. 2021 edition. executive report." https://eepublicdownloads.azureedge.net/clean-documents/sdc-documents/ERAA/ERAA_2021_Executive%20Report.pdf, accessed: 2022-09-26.
- [5] "Methodology for the european resource adequacy assessment." https://www.acer.europa.eu/sites/default/files/documents/Individual%20Decisions_annex/ACER%20Decision%2024-2020%20on%20ERAA%20-%20Annex%20I_1.pdf, accessed: 2022-09-26.
- [6] "European resource adequacy assessment. 2021 edition. annex 3 methodology." https://eepublicdownloads.azureedge.net/clean-documents/sdc-documents/ERAA/ERAA_2021_Annex_3_Methodology.pdf, accessed: 2022-09-26.
- [7] "Methodology for the european resource adequacy assessment." https://eepublicdownloads.azureedge.net/clean-documents/sdc-documents/ERAA/2022/data-for-publication/ERAA2022_Annex_2_Methodology.pdf, accessed: 2023-03-30.
- [8] G. A. Bakirtzis, P. N. Biskas, and V. Chatziathanasiou, "Generation expansion planning by milp considering mid-term scheduling decisions," *Electric Power Systems Research*, vol. 86, pp. 98–112, 2012.
- [9] J. Aghaei, N. Amjadi, A. Baharvandi, and M.-A. Akbari, "Generation and transmission expansion planning: Milp-based probabilistic model," *IEEE Transactions on Power Systems*, vol. 29, no. 4, pp. 1592–1601, 2014.
- [10] E. Gil, I. Aravena, and R. Cárdenas, "Generation capacity expansion planning under hydro uncertainty using stochastic mixed integer programming and scenario reduction," *IEEE Transactions on Power Systems*, vol. 30, no. 4, pp. 1838–1847, 2014.
- [11] S. Backe, M. Ahang, and A. Tomasgard, "Stable stochastic capacity expansion with variable renewables: Comparing moment matching and stratified scenario generation sampling," *Applied Energy*, vol. 302, p. 117538, 2021.
- [12] S. Kamalinia and M. Shahidehpour, "Generation expansion planning in wind-thermal power systems," *IET generation, transmission & distribution*, vol. 4, no. 8, pp. 940–951, 2010.
- [13] S. Jin, S. M. Ryan, J.-P. Watson, and D. L. Woodruff, "Modeling and solving a large-scale generation expansion planning problem under uncertainty," *Energy Systems*, vol. 2, no. 3, pp. 209–242, 2011.
- [14] Y. Feng and S. M. Ryan, "Scenario construction and reduction applied to stochastic power generation expansion planning," *Computers & Operations Research*, vol. 40, no. 1, pp. 9–23, 2013.
- [15] A. P. Hilbers, D. J. Brayshaw, and A. Gandy, "Importance subsampling: improving power system planning under climate-based uncertainty," *Applied Energy*, vol. 251, p. 113114, 2019.
- [16] M. Sun, F. Teng, X. Zhang, G. Strbac, and D. Pudjianto, "Data-driven representative day selection for investment decisions: A cost-oriented approach," *IEEE Transactions on Power Systems*, vol. 34, no. 4, pp. 2925–2936, 2019.
- [17] H. Teichgraber, C. P. Lindenmeyer, N. Baumgärtner, L. Kotzur, D. Stolten, M. Robinius, A. Bardow, and A. R. Brandt, "Extreme events in time series aggregation: A case study for optimal residential energy supply systems," *Applied energy*, vol. 275, p. 115223, 2020.
- [18] J. R. Birge and F. Louveaux, *Introduction to stochastic programming*. Springer Science & Business Media, 2011.
- [19] J. A. Bloom, "Long-range generation planning using decomposition and probabilistic simulation," *IEEE Transactions on Power Apparatus and Systems*, no. 4, pp. 797–802, 1982.
- [20] B. Gorenstin, N. Campodonico, J. Costa, and M. Pereira, "Power system expansion planning under uncertainty," *IEEE transactions on power systems*, vol. 8, no. 1, pp. 129–136, 1993.
- [21] J. Sirikum, A. Techanitisawad, and V. Kachitvichyanukul, "A new efficient ga-benders' decomposition method: For power generation expansion planning with emission controls," *IEEE Transactions on Power Systems*, vol. 22, no. 3, pp. 1092–1100, 2007.
- [22] J. H. Roh, M. Shahidehpour, and L. Wu, "Market-based generation and transmission planning with uncertainties," *IEEE Transactions on Power Systems*, vol. 24, no. 3, pp. 1587–1598, 2009.
- [23] S. Backe, C. Skar, P. C. del Granado, O. Turgut, and A. Tomasgard, "Empire: An open-source model based on multi-horizon programming for energy transition analyses," *SoftwareX*, vol. 17, p. 100877, 2022.
- [24] A. Reiten and M. E. Mikkelsen, "A stochastic capacity expansion model for the european power system—using a distributed progressive hedging approach to handle long-term uncertainty," Master's thesis, NTNU, 2018.
- [25] A. Soares, A. Street, T. Andrade, and J. D. Garcia, "An integrated progressive hedging and benders decomposition with multiple master method to solve the brazilian generation expansion problem," *IEEE Transactions on Power Systems*, vol. 37, no. 5, pp. 4017–4027, 2022.
- [26] H. Park and R. Baldick, "Multi-year stochastic generation capacity expansion planning under environmental energy policy," *Applied energy*, vol. 183, pp. 737–745, 2016.
- [27] S. Gonzato, K. Bruninx, and E. Delarue, "Long term storage in generation expansion planning models with a reduced temporal scope," *Applied Energy*, vol. 298, p. 117168, 2021.
- [28] S. Sen and D. Yakowitz, "A primal-dual subgradient method for time staged capacity expansion planning," *European journal of operational research*, vol. 27, no. 3, pp. 301–312, 1986.
- [29] A. Papavasiliou, S. S. Oren, and R. P. O'Neill, "Reserve requirements for wind power integration: A scenario-based stochastic programming framework," *IEEE Transactions on Power Systems*, vol. 26, no. 4, pp. 2197–2206, 2011.
- [30] I. Aravena and A. Papavasiliou, "Asynchronous lagrangian scenario decomposition," *Mathematical Programming Computation*, pp. 1–50, 2020.
- [31] A. Monticelli, M. Pereira, and S. Granville, "Security-constrained optimal power flow with post-contingency corrective rescheduling," *IEEE Transactions on Power Systems*, vol. 2, no. 1, pp. 175–180, 1987.
- [32] M. Teixeira, H. Pinto, M. Pereira, and M. McCoy, "Developing concurrent processing applications to power system planning and operations," *IEEE Transactions on Power Systems*, vol. 5, no. 2, pp. 659–664, 1990.
- [33] B. H. Kim and R. Baldick, "Coarse-grained distributed optimal power flow," *IEEE Transactions on Power Systems*, vol. 12, no. 2, pp. 932–939, 1997.
- [34] A. G. Bakirtzis and P. N. Biskas, "A decentralized solution to the dc-opf of interconnected power systems," *IEEE Transactions on Power Systems*, vol. 18, no. 3, pp. 1007–1013, 2003.
- [35] A. Papavasiliou, S. S. Oren, and B. Rountree, "Applying high performance computing to transmission-constrained stochastic unit commitment for renewable energy integration," *IEEE Transactions on Power Systems*, vol. 30, no. 3, pp. 1109–1120, 2014.
- [36] M. V. Pereira and L. M. Pinto, "Multi-stage stochastic optimization applied to energy planning," *Mathematical programming*, vol. 52, no. 1-3, pp. 359–375, 1991.
- [37] A. Shapiro, W. Tekaya, J. P. da Costa, and M. P. Soares, "Risk neutral and risk averse stochastic dual dynamic programming method," *European journal of operational research*, vol. 224, no. 2, pp. 375–391, 2013.
- [38] R. J. Pinto, C. T. Borges, and M. E. Maceira, "An efficient parallel algorithm for large scale hydrothermal system operation planning," *IEEE Transactions on Power Systems*, vol. 28, no. 4, pp. 4888–4896, 2013.
- [39] A. Helseth and H. Braaten, "Efficient parallelization of the stochastic dual dynamic programming algorithm applied to hydropower scheduling," *Energies*, vol. 8, no. 12, pp. 14287–14297, 2015.
- [40] F. D. Machado, A. L. Diniz, C. L. Borges, and L. C. Brandão, "Asynchronous parallel stochastic dual dynamic programming applied to hydrothermal generation planning," *Electric Power Systems Research*, vol. 191, p. 106907, 2021.
- [41] D. Ávila, A. Papavasiliou, and N. Löhdorf, "Parallel and distributed computing for stochastic dual dynamic programming," *Computational Management Science*, vol. 19, no. 2, pp. 199–226, 2022.
- [42] D. Avila, A. Papavasiliou, and N. Löhdorf, "Batch learning in stochastic dual dynamic programming," 2021.

- [43] I. Aravena and A. Papavasiliou, "Asynchronous lagrangian scenario decomposition," *Mathematical Programming Computation*, vol. 13, no. 1, pp. 1–50, 2021.
- [44] "European resource adequacy assessment. 2021 edition." <https://www.entsoe.eu/outlooks/eraa/2021>, accessed: 2022-09-26.
- [45] "European resource adequacy assessment. 2021 edition. annex 1 assumptions." https://eepublicdownloads.azureedge.net/clean-documents/sdc-documents/ERAA/ERAA_Annex_1_Assumptions.pdf, accessed: 2022-09-26.
- [46] J. Bezanson, A. Edelman, S. Karpinski, and V. B. Shah, "Julia: A fresh approach to numerical computing," *SIAM review*, vol. 59, no. 1, pp. 65–98, 2017. [Online]. Available: <https://doi.org/10.1137/141000671>
- [47] A. Papavasiliou and S. S. Oren, "Multiarea stochastic unit commitment for high wind penetration in a transmission constrained network," *Operations research*, vol. 61, no. 3, pp. 578–592, 2013.
- [48] A. B. Philpott and Z. Guan, "On the convergence of stochastic dual dynamic programming and related methods," *Operations Research Letters*, vol. 36, no. 4, pp. 450–455, 2008.
- [49] "Subgradient methods," https://web.stanford.edu/class/ee364b/lectures/subgrad_method_slides.pdf, accessed: 2022-09-26.
- [50] D. P. Bertsekas *et al.*, "Dynamic programming and optimal control 3rd edition, volume ii," *Belmont, MA: Athena Scientific*, 2011.



Lazaros Exizidis is an electrical engineer with expertise in power systems. He received the integrated M.Sc. degree in electrical engineering from the Aristotle University of Thessaloniki, Greece (2012) and the Ph.D. degree from the University of Mons, Belgium (2017). Since then, he is an expert with the European Network of Transmission System Operators for Electricity (ENTSO-E). His research interests include optimization methods in power systems, resource adequacy and energy economics.



Daniel Ávila is a PhD student at the Center for Operations Research and Econometrics (CORE), at the Université catholique de Louvain. He received his bachelor and master's degree from the mathematics department of Universidad de los Andes, Bogotá, Colombia. He works under the supervision of Professor Anthony Papavasiliou. His research interests include convex optimization, stochastic programming, renewable energy integration in power systems.



Anthony Papavasiliou (M'06-SM'20) is an Assistant Professor at the department of Electrical and Computer Engineering at the National Technical University of Athens in Greece. He was formerly an Associate Professor and holder of the ENGIE Chair at the Center for Operations Research and Econometrics at the Université catholique de Louvain. He works on operations research, electricity market design, and electric power system operations. He is the recipient of the Francqui Foundation research professorship (2018), the ERC Starting Grant (2019), and the Bodossaki Foundation Distinguished Young Scientist award (2021). He has served as an Associate Editor of *Operations Research* and the *IEEE Transactions on Power Systems*.



Mauricio Junca is an Associate Professor at the Department of Mathematics in Universidad de Los Andes, Bogotá, Colombia. He holds a Ph.D. in Industrial Engineering and Operations Research from UC Berkeley and a MSc. and BSc. in Mathematics from Universidad de Los Andes. His research includes stochastic optimal control, financial mathematics, and stochastic optimization.