This article has been accepted for publication in IEEE Transactions on Power Systems. This is the author's version which has not been fully edited and content may change prior to final publication. Citation information: DOI 10.1109/TPWRS.2026.3725440

Decomposition techniques for model-based bid selection in renewable auctions

Alexandros Visas, *Student Member, IEEE,* Daniel Ávila, Anthony Papavasiliou, *Senior Member, IEEE,* Georg Thomassen, and David Pozo, *Senior Member, IEEE*

***Abstract*—Long-term investments within the European Union have typically relied on zonal price signals, which are unable to capture locational information. This has created a motivation for resorting to locational capacity auctions in order to steer investments to appropriate locations in the grid. The current work provides an algorithm for clearing bids in the context of capacity auctions for renewable investments, while accounting for nodal network constraints at the scale of the full European power system. The corresponding monolithic optimization problem would comprise approximately 760 million variables and 1.5 billion constraints. The proposed framework models bid selection in renewable auctions as a large scale continuous two stage stochastic capacity expansion problem, subject to constraints on the amount of renewable energy that is integrated into the system. The algorithm is designed in order to tackle large scale systems with a particular emphasis on (i) volumes of renewable targets, (ii) nodal resolution, (iii) high temporal resolution and (iv) uncertainty of climate patterns. The developed algorithm relies on a variety of reformulations and decomposition techniques, and it is specifically designed to exploit high performance computing infrastructure.**

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

NOMENCLATURE

Sets and Indices

$F(n)$	Set of lines that originate from node n
$G(n)$	Set of thermal generators at node $n \in N$
$G^{\mathrm{INV}}(n)$	Set of candidate investments at node $n \in N$
$G^{RE}(n)$	Set of existing renewable generators at node $n \in N$
$H^{I}(n)$	Hydro generators at node n, $I \in \{RoR, R, PS\}$[1]
$\mathcal{L}^{AC/DC}$	Set of AC/DC lines
N	Set of nodes
Ω	Set of scenarios
$\mathcal{T}$	Set of time steps of the economic dispatch
$T(n)$	Set of lines that have as a destination node n

Parameters

B_k	Susceptance of line k
$CO2_g$	Cost of CO_2 emissions of technology g (€/MWh)
D_g	Pump capacity of unit g (MW)
η_p/η_g	Pump/turbine efficiency of hydro storage unit (%)
IC_g	Annualized investment cost of technology g (€/MW-year)
$L_k^{\max}$	Maximum capacity of line k (MW)
$L_k^{\min}$	Minimum capacity of line k (MW)
MC_g	Marginal cost of technology g (€/MWh)
$P_g^{\max}$	Maximum generation capacity (MW)
P_g^{RE}	Installed capacity of renewable generator g (MW)
Q_g^I	Turbine capacity of unit g, $I \in \{R, PS\}$ (MW)
TV	Total target volume of net renewable production
V_g^I	Storage capacity of unit g, $I \in \{R, PS\}$ (MWh)
$VOLL$	Value of lost load (€/MWh)
X_g	Capacity expansion limit (MW)

Parameters with Uncertainty

$A_{g,t,\omega}^I$	Hydrological inflow, $I \in \{RoR, R, PS\}$ (MW)
$D_{n,t,\omega}$	Demand at node n, time t, scenario ω (MW)
$TS_{g,t,\omega}$	Normalized time series production profile

Variables

$d_{g,t,\omega}$	Pumped power of unit g (MW)
$\delta_{n,t,\omega}$	Voltage angle at node n
$f_{k,t,\omega}$	Power transferred through line k (MW)
$ls_{n,t,\omega}$	Load shedding (MW)
$p_{g,t,\omega}$	Power produced by thermal generator g (MW)
$p_{g,t,\omega}^{RE}$	Power produced by existing renewable generator g (MW)
$ps_{n,t,\omega}$	Production shedding (MW)
$p_{g,t,\omega}^{xRE}$	Power produced by newly invested renewable generator g (MW)
$q_{g,t,\omega}^I$	Turbined power, $I \in \{RoR, R, PS\}$ (MW)
$s_{g,t,\omega}$	Amount of spillage of unit g (MW)
$v_{g,t,\omega}^I$	State of storage of reservoir g (MWh)
x_g	Capacity invested in technology g (MW)

I. INTRODUCTION

A. Context

Recent geopolitical developments and evolving climate and environmental policies justify the European Union's ambitious drive towards an accelerated integration of renewable energy into the European power grid. This is outlined by a recent communication of the European Commission [1] that proposes a 90% reduction in emissions by 2040. These

Disclaimer: The information and views set out are those of the author(s) and may not in any circumstances be regarded as stating an official position of the European Commission.

Alexandros Visas is with the School of Electrical and Computer Engineering, National Technical University of Athens (NTUA), 15772 Zografou, Greece (email: visas_alexandros@mail.ntua.gr).

Daniel Ávila and Anthony Papavasiliou are with the School of Electrical and Computer Engineering, National Technical University of Athens (NTUA), 15772 Zografou, Greece, and Arktos Analytics, 1040 Etterbeek, Belgium (email: daniel_avila@mail.ntua.gr; papavasiliou@mail.ntua.gr).

Georg Thomassen is with the European Commission, Joint Research Centre (JRC), European Commission, 21027 Ispra, Italy (e-mail: georg.thomassen@ec.europa.eu).

David Pozo is with the European Commission, Joint Research Centre (JRC), 1755 LE Petten, The Netherlands (e-mail: david.pozo-camara@ec.europa.eu).

[1]For ease of exposition, technology-specific quantities, including sets, variables, and parameters, are distinguished by superscripts: run-of-river (RoR), reservoir (R), and pumped storage (PS).

envisioned goals are highly dependent on the large scale deployment of renewable generation resources throughout the grid. Unfortunately, the current European market design does not provide appropriate locational investment incentives [2]. Consequently, this can lead to situations where investments are misplaced and cannot be used to their full potential due to transmission network constraints. Thus, the integration of renewable resources leads to novel challenges in power system planning, since it requires a targeted placement of new plants in locations where the system can integrate these resources.

The European electricity market design is based on a zonal pricing approach. This approach considers an aggregation of the network into large clusters of nodes, referred to as zones. The nodes within a zone share a common energy price when the market clears. In the literature, numerous works have pointed out challenges related to the zonal design. These include the delineation of bidding zones [3], distortions to long-term investment signals [4], [5], and short-term inefficiencies in system operation [6], [7]. From the perspective of investment, zonal designs distort locational signals and thus cloud the ability of the market to stimulate investments at non-congested parts of the network. These distortions may also create arbitrage opportunities between zonal and redispatch prices and thereby further distort long-run investment incentives [5]. Recent studies by the European Commission [7] indicate that the magnitude of the problem can be substantial, with redispatch exceeding 800 TWh and 100 billion € by 2040 under uncoordinated renewable deployment. Instead, locational investment signals alone could reduce system costs by 12-29 billion € per year. A related analysis of renewable auction design finds that redispatch may reach up to 1113 TWh under a zonal auction design that neglects intra-zonal grid constraints, which is more than three times the level observed under a locational auction design [8]. These observations highlight the importance of steering renewable investments toward locations that are compatible with the topology of the network.

Given the tight physical interdependency of national power system operations due to power flow constraints, coordination at the stage of investment will be critical to ensure a least-cost deployment of renewable energy infrastructure. In order to assess the value of coordination and ensure that investments are compatible with the physics that governs the flow of power, it becomes relevant to develop a model that accounts for the entire European continent at nodal resolution.

The volatile and variable supply of renewable energy implies that any generation capacity expansion model should be capable of representing operations at a sufficiently fine temporal resolution to capture renewable supply variability. This further escalates the computational complexity of the task at hand due to strong temporal interactions in addition to the exogenous uncertainty and the broad geographic scope of the required model.

The deployment of renewables will unfold in a highly uncertain environment. Electricity demand depends on climatic conditions, which evolve unpredictably from year to year. Moreover, massive geopolitical uncertainty is introducing major uncertainty in future fuel costs of alternative technologies. Investment in power generation capacity in this uncertain environment can be represented as a two stage stochastic problem [9], where the first stage of decision making is investment and is followed by the operation of the system in a second stage. However, achieving the renewable integration targets envisioned by the European Union motivates the introduction of a renewable integration target, which distorts the standard two stage stochastic programming framework.

Leveraging recent advances in decomposition techniques, we tackle this challenge by proposing an algorithm that clears auctions to select renewable investments on the scale of the entire European system.

B. *Stochastic capacity generation expansion*

Stochastic programming formulations of capacity expansion models have been widely studied in the literature [10]–[21]. This has led to a wide range of solution techniques that are differentiated depending on the specific characteristics of the problem at hand. There are examples of authors who have directly used an optimization solver when the problem is not prohibitively large [14], [15]. There are publications where scenario reduction has been employed in order to reduce the complexity of the problem [16]–[19]. Decomposition methods have also received considerable attention, including Benders decomposition [10]–[12], subgradient-based algorithms [21], and column-generation schemes [13]. A recent alternative framework for large scale stochastic expansion was presented in [22], where a variant of the progressive hedging algorithm is applied to instances with 360 representative days on a system with thousands of nodes.

The decomposition approaches that are most relevant for our setting face three major limitations for our purposes. Firstly, these schemes require the problem to be a standard two stage stochastic problem. Our problem of interest aims at investing while ensuring that the expected net renewable energy integrated into the system reaches a prescribed target. This constraint departs from the typical structure of two stage stochastic problems, as it introduces a coupling constraint between scenarios and thus impedes the direct application of standard decomposition schemes. Secondly, these decomposition algorithms work under the assumption that the subproblems can be solved efficiently. However, the modeling requirements for our problem imply that the subproblems themselves are large scale optimization problems. A third limitation is that the application of high performance computing with these decomposition schemes is typically constrained by the number of scenarios, since a scenario is typically handled by a processor [21]. Our target applications involve just a few tens of scenarios, thus severely limiting our ability to exploit high performance computing clusters with very large numbers of processors.

Some recent developments in decomposition schemes have allowed for a relaxation of temporal coupling constraints by the introduction of so-called budget variables [23], which become first stage decisions. Other recent developments have considered Lagrangian relaxations of scenario coupling constraints and a variety of decomposition algorithms to cope with these scenario coupling constraints [20]. In [20], it is further empirically demonstrated that the approach of budget variables offers a favorable tradeoff in terms of performance and simplicity of implementation. Some recent work has applied the idea of budget variables to extend Benders decomposition to the sectoral and spatial domains in multi-sector capacity expansion models [24].

Leveraging these recent developments, we propose a decomposition scheme for tackling expansion models at EU scale and at nodal resolution. These models lead to large scale optimization problems with thousands of nodes, thousands of time steps, thousands of dispatchable units, and hundreds of storage facilities. The model represents exogenous uncertainty explicitly through scenarios and includes a constraint requiring the net renewable energy absorbed by the system to lie within a target band. The overall solution is tailored to exploit high performance computing infrastructure by allowing a large collection of processors to evaluate subproblems in parallel.

C. Organization and contributions

The current work contributes as follows to the literature. (i) From a methodological perspective, we formulate bid selection in renewable auctions at full-EU scale as a continuous linear problem with a renewable energy integration target band, a nodal representation with more than 4000 nodes, high temporal resolution, and 15 climate scenarios. To the best of our knowledge, this is the first study addressing nodal stochastic generation expansion planning at the scale of the full European power system. (ii) From an algorithmic perspective, we use budget variables, in the spirit of [23], [25], to decompose the stochastic target band constraint and stored energy boundary variables to exactly decompose intertemporal storage constraints while preserving the long-term energy storage behavior of hydro units. The resulting scheme combines an initialization strategy, level set stabilization with dynamic level parametrization, cut consolidation in the spirit of [26], cut sparsification, and problem-specific valid inequalities for both master problems, enabling the solution of an instance the monolithic stochastic equivalent of which would comprise approximately 760 million variables and 1.5 billion constraints. (iii) From an institutional standpoint, this work, carried out in collaboration with the European Commission Joint Research Centre, develops a proof of concept for coordinated renewable auctions as a corrective instrument for weak locational investment signals under the zonal design in Europe.

The paper is organized as follows. Section II presents the mathematical formulation of model-based bid selection in renewable auctions. Section III develops the decomposition scheme for the resulting large scale stochastic problem. Section IV describes the high performance computing implementation enabled by the resulting array of subproblems. Section V presents the case study, and Section VI concludes.

II. Model formulation

The bid selection model in the context of renewable auctions is formulated as a two stage stochastic capacity expansion problem [9]. Nevertheless, the model differs from standard two stage stochastic problems due to the introduction of a second stage constraint that couples scenarios (uncertainty realizations) directly. Within the model, the first stage determines the locations and amounts of investments of renewable assets, while the second stage solves an economic dispatch over a target year. Uncertainty is introduced in the second stage. Each scenario consists of a time series of demand, solar production, wind production, and hydro inflows. The second stage includes a constraint that ensures that the average total net renewable production lies within a target band. The notation used throughout the formulation is summarized in the nomenclature.

A. Investment constraints

The model aims at optimizing the selection of renewable investment bids, specifically for photovoltaic (PV) and wind projects, within the auction framework [8], [27]. There is a maximum amount of investment capacity, which we model with the following constraint:

$$x_g \leq X_g \qquad \forall g \in G^{\mathrm{INV}}(n), n \in N \tag{1}$$

B. Target band constraint

The constraint is modeled as follows:

$$\sum_{\omega \in \Omega} \pi_\omega \sum_{t \in \mathcal{T}} \left(\sum_{\substack{n \in N \\ g \in G^{RE}(n)}} p^{RE}_{g,t,\omega} + \sum_{\substack{n \in N \\ g \in G^{\mathrm{INV}}(n)}} p^{xRE}_{g,t,\omega} - \sum_{n \in N} ps_{n,t,\omega} \right) \in [TV,\ (1+\alpha)\cdot TV] \tag{2}$$

We highlight that this constraint departs from the typical structure of two stage stochastic problems. The introduction of the expected value has the effect of coupling the scenarios directly (as opposed to indirectly in the standard programming literature, i.e. through first stage decisions)[2]. In this expression, π_ω denotes the probability of scenario ω, $p^{RE}_{g,t,\omega}$ and $p^{xRE}_{g,t,\omega}$ denote the production of existing and newly installed renewable units, and $ps_{n,t,\omega}$ denotes production shedding. Subtracting $ps_{n,t,\omega}$ from the expectation in Eq. (2) accounts for production shedding and avoids satisfying the requirement by overbuilding of RES and then shedding their production. Moreover, TV denotes the baseline target level, and α denotes the relative upward deviation allowed above TV. In our application, $\alpha = 5\%$.

C. Photovoltaic and wind generation constraints

Given a scenario $\omega \in \Omega$, we model the production of renewable resources as follows:

$$p^{RE}_{g,t,\omega} = P^{RE}_g \cdot TS_{g,t,\omega} \quad \forall g \in G^{RE}(n), n \in N, t \in \mathcal{T} \tag{3}$$

$$p^{xRE}_{g,t,\omega} = x_g \cdot TS_{g,t,\omega} \quad \forall g \in G^{\mathrm{INV}}(n), n \in N, t \in \mathcal{T} \tag{4}$$

Eq. (3) models the production of existing PV and wind capacity by multiplying installed capacity by a normalized availability profile. Eq. (4) applies the same profile to newly

[2]This implies that existing decomposition techniques, such as Benders decomposition or subgradient-based schemes, are not directly applicable to the problem.

installed capacity x_g. The parameter $TS_{g,t,\omega}$ denotes the availability of renewable generator g at time t under scenario ω.

D. Transmission network

We consider both DC and AC transmission lines within the model. The AC lines are modeled using the B-theta linear approximation of the power flow equations [28]. Given a scenario $\omega \in \Omega$, the B-theta model is formulated using the following constraints:

$$f_{k,t,\omega} - B_k \cdot (\delta_{m,t,\omega} - \delta_{n,t,\omega}) = 0 \tag{5}$$
$$L_k^{min} \leq f_{k,t,\omega} \leq L_k^{max} \tag{6}$$
$$\forall k = (m,n) \in \mathcal{L}^{AC}, t \in \mathcal{T}$$

Eq. (5) links the AC-line flow on line $k = (m,n)$ to the voltage angle difference between its terminal nodes through the susceptance B_k, while Eq. (6) enforces the transmission capacity limits. On the other hand, for a given scenario $\omega \in \Omega$, the DC lines obey the following constraints:

$$L_k^{min} \leq f_{k,t,\omega} \leq L_k^{max} \qquad \forall k \in \mathcal{L}^{DC}, t \in \mathcal{T} \tag{7}$$

Eq. (7) models DC lines as controllable flows that are limited only by their transmission capacities.

E. Thermal generation constraints

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

$$p_{g,t,\omega} \leq P_g^{max} \qquad \forall g \in G(n),\, n \in N,\, t \in \mathcal{T} \tag{8}$$

F. Hydro power

We consider three hydropower generation technologies: run-of-river, reservoir and pumped storage. For a given $\omega \in \Omega$, each one of these technologies is modeled as follows.

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

$$q_{g,t,\omega}^{RoR} = A_{g,t,\omega}^{RoR} \qquad \forall g \in H^{RoR}(n),\, n \in N,\, t \in \mathcal{T} \tag{9}$$

• Reservoir: Water can be stored in the reservoir or released for power generation.

$$v_{g,t,\omega}^{R} = v_{g,t-1,\omega}^{R} + A_{g,t,\omega}^{R} - \frac{q_{g,t,\omega}^{R}}{\eta_g} - s_{g,t,\omega} \tag{10}$$
$$v_{g,t,\omega}^{R} \leq V_g^{R} \tag{11}$$
$$q_{g,t,\omega}^{R} \leq Q_g^{R} \tag{12}$$
$$\forall g \in H^{R}(n),\, n \in N,\, t \in \mathcal{T}$$

Eq. (10) describes the energy storage dynamics of the reservoir. In particular, inflows $A_{g,t,\omega}^{R}$ increase the stored energy, turbine generation $q_{g,t,\omega}^{R}$ depletes the reservoir after accounting for turbine losses through η_g (turbine efficiency), and spillage $s_{g,t,\omega}$ represents water released without generation. Eqs. (11), (12) describe the maximum storage and maximum power production capacities, respectively.

• Pumped storage: We model pumped storage with the following set of constraints:

$$v_{g,t,\omega}^{PS} = v_{g,t-1,\omega}^{PS} + \eta_p \cdot d_{g,t,\omega} - \frac{q_{g,t,\omega}^{PS}}{\eta_g} \tag{13}$$
$$v_{g,t,\omega}^{PS} \leq V_g^{PS} \tag{14}$$
$$q_{g,t,\omega}^{PS} \leq Q_g^{PS} \tag{15}$$
$$d_{g,t,\omega} \leq D_g \tag{16}$$
$$\forall g \in H^{PS}(n),\, n \in N,\, t \in \mathcal{T}$$

Eq. (13) describes the energy balance of the pumped-storage reservoir, where η_p and η_g denote the pumping and turbine efficiencies, respectively. Thus, stored energy increases by less than the electricity consumed during pumping and decreases by more than the electricity delivered during generation.[3] Eq. (14) bounds the maximum storage, while Eqs. (15), (16) bound the turbine output and pumping power respectively.

G. Nodal energy balance

For each scenario $\omega \in \Omega$, nodal energy balance is enforced as follows:

$$\begin{aligned}
&\sum_{g \in G(n)} p_{g,t,\omega} + \sum_{\substack{I \in \{RoR,R,PS\} \\ g \in H^I(n)}} q_{g,t,\omega}^{I} \\
&+ \sum_{g \in G^{RE}(n)} p_{g,t,\omega}^{RE} + \sum_{g \in G^{\mathrm{INV}}(n)} p_{g,t,\omega}^{xRE} + \sum_{k \in T(n)} f_{k,t,\omega} \\
&- \sum_{k \in F(n)} f_{k,t,\omega} - \sum_{g \in H^{PS}(n)} d_{g,t,\omega} + ls_{n,t,\omega} - ps_{n,t,\omega} \\
&= D_{n,t,\omega} \quad \forall n \in N,\, t \in \mathcal{T}
\end{aligned} \tag{17}$$

Eq. (17) balances, at each node and time period, thermal, hydro, and renewable generation, net imports, pumping consumption, load shedding, and production shedding. The variable $ps_{n,t,\omega}$ represents production shedding and absorbs excess renewable output when this output cannot be accommodated by the system.

The load shedding term in Eq. (17) is bounded by nodal demand for each scenario $\omega \in \Omega$:

$$ls_{n,t,\omega} \leq D_{n,t,\omega} \qquad \forall n \in N,\, t \in \mathcal{T} \tag{18}$$

H. Objective function

The term (19) is the first stage cost and consists of the investment cost.

$$\sum_{n \in N, g \in G^{\mathrm{INV}}(n)} IC_g \cdot x_g \tag{19}$$

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

$$\sum_{\substack{n \in N,\, g \in G(n) \\ t \in \mathcal{T}}} (MC_g + CO2_g) \cdot p_{g,t,\omega} \tag{20}$$

[3]Short-term storage technologies, such as battery units, can be modeled analogously by introducing upper bounds on the storage level and on the charging and discharging variables, as in Eqs. (14)-(16), together with an intertemporal balance equation with charging and discharging efficiencies, as in Eq. (13).

$$\sum_{\substack{n \in N \\ t \in \mathcal{T}}} VOLL \cdot ls_{n,t,\omega} \tag{21}$$

The cost of producing $p_{g,t,\omega}$ units of power is described by term (20). Term (21) assigns the VOLL penalty to the nodal load shedding variable introduced in (18).

I. Problem formulation

Putting these elements together leads to the problem of model-based bid selection in renewable auctions (MRA), which is described as follows.

$$\begin{aligned} \text{(MRA)} \quad & \min_{x, \{y_\omega\}_{\omega \in \Omega}} (19) + \sum_{\omega \in \Omega} \pi_\omega \Big[(20) + (21) \Big] \\ & \text{s.t.} \quad (1),\ (2) \\ & \qquad (3) - (18) \quad \forall \omega \in \Omega \end{aligned}$$

For compactness, we denote by y_ω[4] the collection of second stage operational variables under scenario ω. The constraints in (MRA) are separated into three sets. The first set contains the first stage constraints, represented by (1), which do not depend on $\omega \in \Omega$. The second set corresponds to the scenario-coupling constraint, represented by (2). The third set contains the second stage operational constraints, represented by (3)-(18), which depend on uncertainty and are imposed separately for each scenario $\omega \in \Omega$.

III. Decomposition Strategy

As discussed in section II, the problem of interest consists of three main blocks of constraints. In subsection III-A, budget variables [23] decouple the scenario-coupling target band constraint, yielding a standard two stage stochastic problem. Subsection III-B applies the L-shaped method to handle the linkage between the first and second stage. Because the economic dispatch within each scenario is too large to solve in a single shot, subsection III-C introduces a temporal decomposition based on budget variables [25] and intertemporal subproblems [21]. Subsection III-D describes a fixed parameter level set stabilization scheme. Subsection III-E presents several enhancements, including a dynamic level parameter, cut consolidation, cut sparsification, and valid inequalities, whose detailed derivations are provided in the electronic supplement [29].

A. Decomposition of the target band constraint

The typical structure of two stage stochastic problems follows the logic that once the first stage decision is determined, uncertainty materializes, and one can react to each scenario without taking into account other scenarios. The situation with the scenario coupling constraint of Eq. (2) that is present in this problem introduces a more complex setting. Once the first stage decision is determined, one cannot consider each scenario independently of every other scenario, since there is a constraint that couples the decisions taken across different scenarios.

To tackle the coupling between scenarios, we resort to the methodology developed in [25]. For each scenario $\omega \in \Omega$ we introduce a so-called budget variable q_ω, representing the net renewable production budget allocated to scenario ω, and reformulate the (MRA) problem as:

$$\begin{aligned} \text{(B-MRA)} \quad & \min_{x, q, \{y_\omega\}_{\omega \in \Omega}} (19) + \sum_{\omega \in \Omega} \pi_\omega \Big[(20) + (21) \Big] \\ & \text{s.t.} \ (1) \\ & \sum_{\omega \in \Omega} \pi_\omega q_\omega = TV \qquad (22) \\ & \sum_{t \in \mathcal{T}} \Bigg(\sum_{\substack{n \in N \\ g \in G^{RE}(n)}} p^{RE}_{g,t,\omega} + \sum_{\substack{n \in N \\ g \in G^{\text{INV}}(n)}} p^{xRE}_{g,t,\omega} - \sum_{n \in N} ps_{n,t,\omega} \Bigg) \\ & \qquad \geq q_\omega \quad \forall \omega \in \Omega \qquad (23) \\ & \sum_{t \in \mathcal{T}} \Bigg(\sum_{\substack{n \in N \\ g \in G^{RE}(n)}} p^{RE}_{g,t,\omega} + \sum_{\substack{n \in N \\ g \in G^{\text{INV}}(n)}} p^{xRE}_{g,t,\omega} - \sum_{n \in N} ps_{n,t,\omega} \Bigg) \\ & \qquad \leq (1+\alpha) \cdot q_\omega \quad \forall \omega \in \Omega \qquad (24) \\ & (3) - (18) \quad \forall \omega \in \Omega \end{aligned}$$

In reformulation (B-MRA), the scenario-coupling target band constraint (2) is replaced by the (master) constraint (22) and the scenario-wise operational constraints (23)-(24). The auxiliary (master) variable q_ω allocates a net renewable production budget to scenario ω. Eq. (22) ensures that these budgets recover the baseline target TV in expectation, while Eqs. (23)-(24) require the realized net renewable production in each scenario to lie between q_ω and $(1+\alpha)q_\omega$. Hence, the coupling across scenarios is moved to the master problem, and the operational constraints become scenario-wise separable. The exactness of this reformulation follows from Theorem 1 of [25], which establishes the equivalence of the budget variable reformulation.

B. Decomposition at scenario level

The L-shaped method allows us to break down a standard two stage stochastic problem [9]. We decouple the first and second stage into a master and subproblem structure. We then create a subproblem for each uncertainty realization for the second stage.

We proceed by considering the reformulation of (B-MRA). For each $\omega \in \Omega$, we define the value function $\mathcal{V}_\omega(x, q_\omega)$ as the minimum second stage operating cost for given first stage decisions x and q_ω. For fixed master values $\bar{x}$ and $\bar{q}_\omega$, the corresponding optimization problem reads:

$$\begin{aligned} \mathcal{V}_\omega(\bar{x}, \bar{q}_\omega) = & \min_{y_\omega} . (20) + (21) \\ & \text{s.t.} \ (23) - (24) \\ & \quad (3) - (18) \end{aligned}$$

[4]Here, y_ω collects all second stage operational variables under scenario ω. When temporal decomposition is introduced, $y_{k,\omega}$ denotes the restriction of these operational variables to temporal segment k under scenario ω.

This article has been accepted for publication in IEEE Transactions on Power Systems. This is the author's version which has not been fully edited and content may change prior to final publication. Citation information: DOI 10.1109/TPWRS.2026.3725440

Using $\mathcal{V}_\omega(x, q_\omega)$, we re-write the (B-MRA) problem as:

$$\begin{aligned} \min_{x,q}.\quad & (19) + \mathbb{E}_\omega\Big[\mathcal{V}_\omega(x, q_\omega)\Big] \\ \text{s.t.}\quad & (1),\ (22) \end{aligned}$$

Standard results from the two stage stochastic programming literature imply that the function $\mathbb{E}_\omega\Big[\mathcal{V}_\omega(x, q_\omega)\Big]$ is convex and piecewise linear on its domain in the first stage decisions x and $q = \{q_\omega\}_{\omega\in\Omega}$ [9]. The L-shaped method under-approximates this function by means of supporting hyperplanes, commonly referred to as optimality cuts, which are derived from the subgradients of the value functions $\mathcal{V}_\omega(x, q_\omega)$ [9]. In the present formulation, we adopt a multi-cut variant of the L-shaped method and generate one optimality cut per scenario at each iteration. As discussed in [30], such multicut variants typically yield a tighter master approximation, although they increase the size of the master problem.

A related consideration concerns the domain of $\mathbb{E}_\omega\Big[\mathcal{V}_\omega(x, q_\omega)\Big]$. When the second stage problem is infeasible for a given first stage decision, the standard remedy is to add feasibility cuts derived from rays of unboundedness [9]. In our approach, we instead introduce slack variables, highly penalized in the objective, in order to arrive at a formulation with complete recourse [9], so that $\mathcal{V}_\omega(x, q_\omega)$ is finite for every pair of first stage decisions (x, q_ω). This choice is motivated by the fact that extracting rays of unboundedness requires the simplex algorithm, which is significantly slower than the barrier algorithm for the problem at hand. The resulting algorithm therefore relies exclusively on optimality cuts.

For each $\omega \in \Omega$ and iteration ν, let c^ν_ω denote the cut generated under scenario ω at iteration ν. Since one cut is generated per scenario at each iteration, the master problem at iteration j contains the cuts $\{c^\nu_\omega(x, q_\omega)\}_{\nu=1}^{j-1}$ for all $\omega \in \Omega$. The previous problem can then be approximated as follows:

$$\begin{aligned} \min_{x,q,\theta}.\quad & (19) + \mathbb{E}_\omega[\theta_\omega] \\ \text{s.t.}\quad & (1),\ (22) \\ & \theta_\omega \geq c^\nu_\omega(x, q_\omega) \quad \forall \nu \in \{1, \dots, j-1\},\ \omega \in \Omega. \end{aligned}$$

This problem is commonly referred to as the master problem of the L-shaped method. Within iteration j, the master problem is first solved, yielding the trial point $(\bar{x}^j, \bar{q}^j)$. Each subproblem $\mathcal{V}_\omega(\bar{x}^j, \bar{q}^j_\omega)$ is then solved, and a subgradient of $\mathcal{V}_\omega$ at $(\bar{x}^j, \bar{q}^j_\omega)$ is used to compute an optimality cut c^j_ω, which is added to the master problem. Throughout this section and in the remainder of the paper, barred quantities denote fixed values passed from the master problem to the subproblems. As the cuts accumulate, the master approximation of $\mathbb{E}_\omega\Big[\mathcal{V}_\omega(x, q_\omega)\Big]$ is progressively refined, and finite convergence is guaranteed [9].

An upper bound is computed at the current trial point by adding the first stage investment cost to the expected second stage operating cost evaluated at that point. A lower bound is given by the optimal value of the master problem. The resulting optimality gap serves as a stopping criterion.

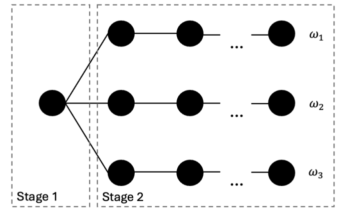

Figure 1. Temporal decomposition of the second stage. The second stage is partitioned into several consecutive chronological segments.

C. Decomposition of intertemporal constraints

Due to the scale of the problem, solving each scenario subproblem $\mathcal{V}_\omega(x, q_\omega)$ in a single shot is computationally intractable. Loading the full subproblem into memory is prohibitive, and even in simplified instances where this is possible, the solution time remains unacceptably large. We therefore seek an additional level of decomposition within each scenario subproblem. To this end, we combine the chronological decomposition of [21] with stored energy boundary variables that link adjacent temporal segments and preserve intertemporal storage dynamics.

We partition the second stage horizon $\{1, \dots, T\}$ into the set of K consecutive temporal segments $\mathcal{K} = \{1, \dots, K\}$, namely $\{1, \dots, t_1\}, \{t_1 + 1, \dots, t_2\}, \dots, \{t_{K-1} + 1, \dots, T\}$, where $t_0 = 0$ and $t_K = T$, as illustrated in Figure 1. For each temporal segment k and scenario ω, and for fixed master values $\bar{x}$, $\bar{q}_{k,\omega}$, $\bar{v}_{t_{k-1},\omega}$, and $\bar{v}_{t_k,\omega}$, we define the following subproblem:

$$\mathcal{V}_{k,\omega}(\bar{x}, \bar{q}_{k,\omega}, \bar{v}_{t_{k-1},\omega}, \bar{v}_{t_k,\omega}) =$$

$$\begin{aligned} \min_{y_{k,\omega}}.\quad & \Big[(20) + (21)\Big]_{t=t_{k-1}+1}^{t_k} \\ \text{s.t.}\quad & \Big[(3) - (18)\Big]_{t=t_{k-1}+1}^{t_k} \end{aligned}$$

$$\sum_{t=t_{k-1}+1}^{t_k}\Bigg(\sum_{\substack{n\in N\\ g\in G^{RE}(n)}} p^{RE}_{g,t,\omega} + \sum_{\substack{n\in N\\ g\in G^{\text{INV}}(n)}} p^{xRE}_{g,t,\omega} - \sum_{n\in N} ps_{n,t,\omega}\Bigg) \geq \bar{q}_{k,\omega} \tag{25}$$

$$\sum_{t=t_{k-1}+1}^{t_k}\Bigg(\sum_{\substack{n\in N\\ g\in G^{RE}(n)}} p^{RE}_{g,t,\omega} + \sum_{\substack{n\in N\\ g\in G^{\text{INV}}(n)}} p^{xRE}_{g,t,\omega} - \sum_{n\in N} ps_{n,t,\omega}\Bigg) \leq (1+\alpha)\bar{q}_{k,\omega} \tag{26}$$

$$v^I_{g,t_{k-1},\omega} = \bar{v}^I_{g,t_{k-1},\omega} \tag{27}$$

$$v^I_{g,t_k,\omega} = \bar{v}^I_{g,t_k,\omega} \tag{28}$$

$$\forall I \in \{R, PS\},\ g \in H^I(n),\ n \in N$$

In the preceding problem, the bracket notation indicates that the corresponding terms are restricted to the time interval $t \in \{t_{k-1} + 1, \dots, t_k\}$. Eqs. (25)-(26) assign to temporal segment k under scenario ω a net renewable production budget $\bar{q}_{k,\omega}$ and require the realized net renewable production over that segment to lie between $\bar{q}_{k,\omega}$ and $(1 + \alpha)\bar{q}_{k,\omega}$, while Eqs. (27)-(28) prescribe the hydro storage level at the beginning and at the end of the segment. In circular indexing approximations such as those adopted in [23], [25], the storage level at the end of each subperiod is set equal to the storage level at its beginning. As a result, hydro storage is represented separately within each subperiod and long term storage dynamics are not represented accurately. Under the proposed reformulation, adjacent temporal segments are linked through common storage boundary variables[5], so that the hydro trajectory remains chronologically consistent over the full horizon. This exact treatment of hydro storage allows problem (B-MRA) to be reformulated as follows.

$$\begin{aligned} \min_{x,q,v}\quad & (19) + \mathbb{E}_\omega\Big[\sum_{k\in\mathcal{K}} \mathcal{V}_{k,\omega}(x, q_{k,\omega}, v_{t_{k-1},\omega}, v_{t_k,\omega})\Big] \\ \text{s.t.}\quad & (1) \\ & \Big[(11), (14)\Big]_{\substack{k\in\mathcal{K}\\ \omega\in\Omega}} \\ & \mathbb{E}_\omega\Big[\sum_{k\in\mathcal{K}} q_{k,\omega}\Big] = TV \qquad (29) \end{aligned}$$

In the temporally decomposed reformulation, the master problem determines the investment vector x, the segment-scenario net renewable production budgets $q_{k,\omega}$, and the stored energy boundary variables $v_{t_k,\omega}$. Eq. (29) ensures that the net renewable production budgets recover the baseline target TV in expectation, while Eqs. (25)-(26) enforce the segment-level tolerance around each budget. Constraints (11) and (14) bound the feasible values of the storage boundary variables.

Following the procedure described in subsection III-B, we apply the L-Shaped method to this temporally decomposed reformulation. The segment-scenario value functions $\mathcal{V}_{k,\omega}$ are treated analogously to the scenario value functions of subsection III-B, but now with respect to the first stage decisions x, $q_{k,\omega}$, $v_{t_{k-1},\omega}$, and $v_{t_k,\omega}$. At each iteration j, one cut $c^j_{k,\omega}$ is generated for every pair (k, ω), and the resulting master problem reads:

$$\begin{aligned} \text{(M)}\quad \min_{x,q,v,\theta}\quad & (19) + \mathbb{E}_\omega\Big[\sum_{k\in\mathcal{K}} \theta_{k,\omega}\Big] \\ \text{s.t.}\quad & (1),\ (29) \\ & \Big[(11),\ (14)\Big]_{\substack{k\in\mathcal{K}\\ \omega\in\Omega}} \\ & \theta_{k,\omega} \geq c^\nu_{k,\omega}(x, q_{k,\omega}, v_{t_{k-1},\omega}, v_{t_k,\omega}) \\ & \quad \forall \nu \in \{1,\dots,j-1\},\ k\in\mathcal{K},\ \omega\in\Omega \end{aligned}$$

As in subsection III-B, the master problem (M) provides a lower bound at each iteration, and the segment-scenario subproblems are evaluated at the resulting trial point to compute an upper bound and generate new cuts. Complete recourse ensures feasibility and thus there is no need for feasibility cuts. Because $|\mathcal{K}| \cdot |\Omega|$ new cuts are added per iteration, (M) contains a total of $(j-1)\cdot|\mathcal{K}|\cdot|\Omega|$ cuts before iteration j, provided that no previously generated cuts are dropped.

[5]For short duration storage technologies, such as batteries, introducing storage boundary variables in the master problems can be avoided. Instead, their stored energy can be approximated within each temporal segment, for example by imposing a cyclic energy condition or by prescribing a representative charge-discharge pattern. Such approximations do not increase the number of first stage variables and are generally reasonable for batteries, since their limited energy capacity usually prevents them from materially affecting system states across long planning horizons.

The advantage of this decomposition is that each segment-scenario subproblem $\mathcal{V}_{k,\omega}$ is significantly smaller and therefore faster to solve. On the other hand, the number of subproblems and the number of first stage decision variables both increase, and the master problem grows faster since the number of cuts per iteration now scales with $|\mathcal{K}|\cdot|\Omega|$ rather than $|\Omega|$ alone.

D. Level set stabilization scheme

Stabilization aims to address well-known drawbacks of cutting plane methods. In early iterations, the approximation may contain little information about the neighborhood of an optimal solution, and the algorithm may spend many iterations producing trial points that explore extreme points of the feasible region [9].

A broad literature has developed around stabilizing Benders algorithms. Trust region, proximal bundle, and level set methods have all been proposed to mitigate oscillatory behavior and improve the quality of trial points in early iterations [30], [31]. Proximal bundle methods penalize deviations from a stability center through a quadratic term in the objective, whereas trust region methods impose an equivalent restriction by constraining the next iterate to a neighborhood of that center [31]–[33]. Level set methods [34] instead project the stability center onto a level set of the cutting plane model, with the level defined through an interpolation between lower and upper bounds [31]. In the energy systems literature, [23] studies several level set variants for capacity expansion, and [32] compares trust region and bundle stabilization techniques for energy planning under climate uncertainty. Related level methods have also been used in Lagrangian decomposition, for example to compute locational convex hull prices [35] and to address stochastic expansion planning with endogenous adequacy targets [20].

Among the available stabilization schemes, preliminary experiments on smaller instances of our problem indicated that level set stabilization outperforms the proximal bundle method for our application. We therefore adopt level set stabilization and first describe a fixed parameter variant, where the level is determined by a prescribed value of λ. Let $z = (x, q, v)$ denote the vector of first stage decisions and let $\widehat{z}$ be the current stability center, that is, the incumbent first stage point associated with the best upper bound found so far. At iteration j, let LB^j be the lower bound provided by the master problem (M), let $(UB^{j-1})^* = \min_{0\le\nu<j} UB^\nu$ denote the best upper bound available prior to iteration j, and let $0 < \lambda < 1$ be the fixed level parameter, so that $\lambda^j = \lambda$ for all j. Defining the level $\ell^j(\lambda^j) = \lambda^j \cdot LB^j + (1-\lambda^j)\cdot(UB^{j-1})^*$ as a convex combination of the current lower and upper bounds leads to:

$$\begin{aligned} \text{(LS-M)}\quad \min_{x,q,v,\theta}\quad & \|z - \widehat{z}\|_2^2 \\ \text{s.t.}\quad & (1),\ (29) \\ & \Big[(11),\ (14)\Big]_{\substack{k\in\mathcal{K}\\ \omega\in\Omega}} \\ & \theta_{k,\omega} \geq c^\nu_{k,\omega}(x, q_{k,\omega}, v_{t_{k-1},\omega}, v_{t_k,\omega}) \\ & \quad \forall \nu\in\{1,\dots,j-1\},\ k\in\mathcal{K},\ \omega\in\Omega \\ & (19) + \mathbb{E}_\omega\Big[\textstyle\sum_{k\in\mathcal{K}} \theta_{k,\omega}\Big] \leq \ell^j(\lambda^j) \qquad (30) \end{aligned}$$

This article has been accepted for publication in IEEE Transactions on Power Systems. This is the author's version which has not been fully edited and content may change prior to final publication. Citation information: DOI 10.1109/TPWRS.2026.3725440

Problem (LS-M) preserves the feasible region described by the cutting plane model and selects, among all points satisfying (30), the one closest to the stability center $\widehat{z}$. We denote the optimal solution of this problem by $\bar{z}^j$. At each iteration, $\bar{z}^j$ is the trial point returned by the level set master problem and is passed to the subproblems for evaluation. The upper bound UB^j is computed by adding the investment cost at $\bar{x}^j$ to the scenario-weighted sum of the segment subproblem values evaluated at $\bar{z}^j$. The stability center $\widehat{z}$ is updated only when $\bar{z}^j$ yields an improved incumbent upper bound. Since (LS-M) includes the same multi-cuts as (M), its size increases analogously, provided that no cuts are dropped.

The objective used in (LS-M) is not unique. In [23], the quadratic choice $\|z-\widehat{z}\|_2^2$ is compared with an unregularized level set variant, in which any point satisfying the level set and cutting-plane constraints can be selected as the next trial point. In our setting, however, this variant would not be able to exploit the initialization strategy described in Algorithm S1 of the electronic supplement [29], because it does not contain any term measuring the distance from the initialized stability center $\widehat{z}^0$. By contrast, distance-based objectives such as $\|z-\widehat{z}\|_1$, $\|z-\widehat{z}\|_2^2$, or $\|z-\widehat{z}\|_\infty$ reflect the magnitude of the deviations from the stability center and therefore guide the next iterate toward promising trial points [36]. In our tests, the ℓ_2-norm combined with the initialization described in Algorithm S1 of the electronic supplement [29] outperforms the corresponding ℓ_1-norm and ℓ_∞-norm alternatives, as well as the unregularized level set variant proposed in [23].

The resulting fixed parameter level set decomposition scheme is presented as Algorithm S1 of the electronic supplement [29]. The initial stability center $\widehat{z}^0$ is obtained by solving one deterministic wait-and-see problem per scenario and averaging the resulting first stage decisions. This provides a meaningful starting point and reduces the number of iterations required for convergence.

E. Enhancements of the stabilization algorithm

Several enhancements are introduced on top of the fixed parameter level set scheme of Algorithm S1 of the electronic supplement [29]. Their detailed derivations are provided in the electronic supplement [29]. The key ideas are summarized below.

- Dynamic level parameter: Motivated by [37], we replace the fixed λ with a coefficient $\lambda^j \in [\lambda_{\min}, \lambda_{\max}]$ that evolves over iterations. Unlike [37], our progress ratios and update steps are tailored to the bound evolution observed in our application, where the incumbent upper bound typically improves early while later convergence is mainly driven by lower bound improvements. A quadratic step adjusts the parameter when both bounds stall, a milder linear correction is applied when only one bound stalls, and the parameter is left unchanged when progress is satisfactory. The detailed update procedure is presented as Routine S2 of the electronic supplement [29], while the resulting level set algorithm with dynamic parametrization is summarized in Algorithm S3 of the same supplement.
- Cut consolidation: Following the consolidation technique of [26], inactive cuts are periodically aggregated. In our temporal multi-cut setting, cuts sharing the same temporal segment and iteration are collected into a single group and, when a sufficient fraction has been inactive for enough consecutive iterations, they are replaced by a single aggregated cut. Related strategies include adaptive multicut aggregation [38] and cut selection based on violations [39]. Because cut consolidation weakens the cutting-plane approximation, the lower bound computed by the LP master problem may temporarily decrease. To avoid this, we impose in both master problems the lower bound preservation constraint reported in the electronic supplement [29], which requires the LP master objective expression to remain at least as large as the best lower bound obtained in previous iterations.
- Cut sparsification: In the proposed algorithm, during cut construction, coefficients with absolute value at most 10^{-6} are set to zero. This removes numerically negligible terms, yields sparser cuts, and improves numerical stability.
- Valid inequalities: To improve the quality of early trial points, we strengthen the master problems with valid inequalities, following the general idea in [30], [32], [40], while deriving the specific inequalities from the structure of our subproblems. These include RES production capacity bounds that anticipate the penalty cost when the renewable target exceeds available capacity, intertemporal bounds on pumped storage that account for physical limits on turbine output and pumping power, and reservoir inflow bounds that account for the nonnegativity of spillage. The derivation of these inequalities and the resulting strengthened master problems, denoted (VC-M) and (VC-LS-M), are presented in the electronic supplement [29].

Algorithm 1 Proposed Decomposition Algorithm

INPUT: λ^1, $j_{\max}$, ϵ_{tol}, $\tau_{\max}$, γ, $r_{\min}$, $r_{\max}$, $\lambda_{\min}$, $\lambda_{\max}$, N_c, R_c[6].
INITIALIZATION: Solve the deterministic wait-and-see problem for each scenario and define $\widehat{z}^0 = (\widehat{x}^0, \widehat{q}^0, \widehat{v}^0)$ as the average of the resulting x, q, and v decisions. Set $\widehat{z} \leftarrow \widehat{z}^0$, initialize $(UB^0)^*$ with a sufficiently large value, set $j \leftarrow 1$, and $\epsilon^0 \leftarrow +\infty$.
while $j \le j_{\max}$, $\epsilon^{j-1} > \epsilon_{\text{tol}}$, and elapsed time $< \tau_{\max}$ **do**
  Solve (VC-M) and set LB^j equal to its objective value.
  If $j \ge 2$, update λ^j with Routine S2 of [29].
  Compute $\ell^j(\lambda^j) = \lambda^j \cdot LB^j + (1-\lambda^j) \cdot (UB^{j-1})^*$.
  Solve (VC-LS-M) to obtain $\bar{z}^j = (\bar{x}^j, \bar{q}^j, \bar{v}^j)$.
  for $k \in \mathcal{K}$, $\omega \in \Omega$ **do**
    Solve $\mathcal{V}_{k,\omega}(\bar{z}^j)$ and generate cut $c^j_{k,\omega}$.
  end for
  Compute UB^j, set $(UB^j)^* = \min\{(UB^{j-1})^*, UB^j\}$.
  If $(UB^j)^* < (UB^{j-1})^*$, set $\widehat{z} \leftarrow \bar{z}^j$.
  Compute $\epsilon^j = \big((UB^j)^* - LB^j\big)/LB^j$.
  Add the cuts $c^j_{k,\omega}$ to (VC-M) and (VC-LS-M).
  Consolidate inactive cuts as described in [29].
  Set $j \leftarrow j+1$.
end while

[6]All input parameters of Algorithm 1 are described in the electronic supplement [29].

This article has been accepted for publication in IEEE Transactions on Power Systems. This is the author's version which has not been fully edited and content may change prior to final publication. Citation information: DOI 10.1109/TPWRS.2026.3725440

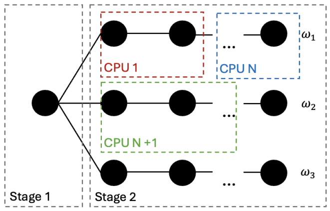

Figure 2. Parallel computing strategy. The set of second stage subproblems is distributed among the available CPUs.

Algorithm 1 summarizes the proposed decomposition algorithm with the enhancements described above.

IV. Parallelization scheme

As discussed in section III-C, the proposed approach increases the number of subproblems relative to standard two stage stochastic programming decomposition techniques. This section describes the parallelization scheme used to exploit high performance computing (HPC) infrastructure. In standard scenario decomposition, the available parallelism is typically limited by the number of scenarios. With temporal splitting, the decomposition creates a larger collection of independent segment-scenario subproblems, which allows a finer distribution of work across the available workers.

In our implementation, the subproblems are assigned to workers through a static allocation fixed before the iterative procedure starts. As discussed in [39], this avoids the communication overhead associated with relocating subproblems during the algorithm.

Within the developed scheme, the set of subproblems is

$$\left\{\mathcal{V}_{k,\omega}(\cdot,\cdot,\cdot,\cdot)\right\}_{k\in\mathcal{K},\omega\in\Omega}.$$

All subproblems in this set are independent, unlike temporal decomposition techniques such as [21], where subproblems remain linked. Consequently, once a master solution has been computed, all segment-scenario subproblems can be evaluated in parallel. Assigning one worker to each subproblem allows the implementation to exploit up to $|\mathcal{K}|\cdot|\Omega|$ worker CPUs, resulting in the parallel computing structure presented in Figure 2. When fewer workers are available, the static allocation assigns several segment-scenario subproblems to each worker.

V. Case study

This section presents the results of the proposed decomposition strategy for a realistic all-EU instance of the power network. The dataset spans 34 countries[7] and an overview of its main characteristics is provided in Table I.

[7]Compared with the entries listed on the ENTSO-E members page, the dataset does not include Cyprus, Iceland, Northern Ireland, or Ukraine. It additionally includes Great Britain and Kosovo, which are modeled in the dataset but are not taken into account in the target volume calculation.

[8]The term solar hsat denotes solar units with horizontal single-axis tracking.

Table I
Overview of the dataset

	Characteristics	Data source
Geography	- Full European system 34 countries	
Transmission network	- Nodal model - 4054 buses, - 5704 AC lines, 36 DC lines	PyPSA [41]
Thermal generation	- Unit-based model - 1491 generation units	PyPSA
Hydrological generation	- Unit-based model - 3 technologies: run-of-river, reservoir, pumped storage - 595 reservoirs, 142 pumped-storage units, 612 run-of-river units	PyPSA
Solar & wind generation	- 3942 solar units, 4050 wind units	PyPSA
Candidate RES investments	- 471 solar, 453 solar hsat,[8] 396 onshore wind, 94 offshore wind projects	PyPSA
Uncertainty	- 15 scenarios - Each scenario consists of a yearly demand, inflow, solar and wind profiles	PyPSA ERAA [42]
Time horizon	- One operational year with 6 time periods per day	

A. Algorithmic setup

The decomposition algorithm is implemented in Julia v1.12.4 using the JuMP v1.30.0 modeling framework, and all linear and quadratic continuous problems are solved with Gurobi 10.0.3. The computational experiments are carried out on the ARIS HPC facility, operated by the National Infrastructures for Research and Technology S.A. of Greece (GRNET S.A.). ARIS comprises 48 thin computational nodes (512 GB RAM each) and 16 fat computational nodes (1 TB RAM each), all equipped with 128 AMD EPYC 7763 cores per node. For the full-EU instance, we allocate 4 fat nodes, using 27 cores and 650 GB of RAM per node, for a total of 108 cores and 2600 GB of RAM.

Based on empirical experiments, we opt to partition each year into 14 temporal segments per scenario. This yields $|\Omega|\cdot 14 = 210$ segment-scenario subproblems in total and provides a favorable tradeoff between the solution time of the subproblems and the size of the two master problems. Additional sensitivity results supporting this choice are reported in the electronic supplement [29]. The collection of subproblems is distributed across the available cores, thus leading to 1-2 subproblems per core. Each subproblem is solved using one Gurobi thread with the barrier algorithm, which we have empirically observed to be faster than the simplex (or dual simplex) algorithm for our application. Empirical evidence from previous research (e.g. [20]) indicates that for these types of problems, enabling the multi-threading option in Gurobi severely increases the computational usage of resources in the cluster without providing a substantial decrease in run time.

This article has been accepted for publication in IEEE Transactions on Power Systems. This is the author's version which has not been fully edited and content may change prior to final publication. Citation information: DOI 10.1109/TPWRS.2026.3725440

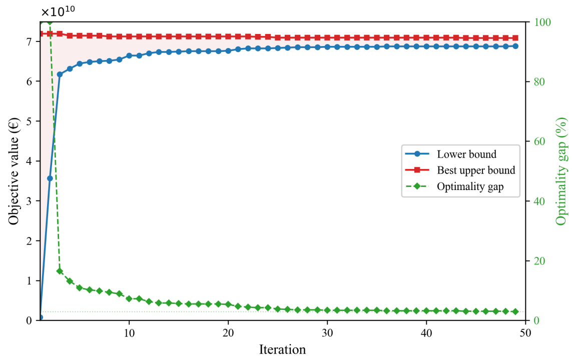

(a) Convergence over iterations.

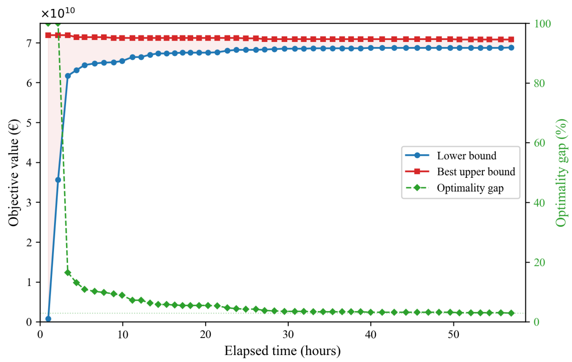

(b) Convergence over elapsed time.

Figure 3. Convergence evolution of the full-EU instance. The upper and lower panels report convergence over iterations and elapsed time, respectively. In both panels, the lower and best upper bounds are shown on the left axis, while the optimality gap is shown on the right axis.

Furthermore, we set the crossover option to 0 to decrease the run time per iteration.

B. Initialization strategy

As described in Algorithm 1, the decomposition scheme is initialized at the stability center $\widehat{z}^0$, obtained from the average wait-and-see solution. For smaller instances, this requires solving one deterministic problem per scenario and averaging the first stage decisions. However, for the full European instance each monolithic single scenario problem comprises approximately 100 million constraints and 51 million variables, rendering a direct solve intractable.

To circumvent this limitation, we partition the European network into four regional subsystems (Central-Western Europe, Northern Europe, Eastern Europe, and Southern Europe) and solve a deterministic problem for each region independently. The resulting regional solutions are then combined to form the initial stability center for the full-EU instance. This regional initialization provides the algorithm with a reasonable starting point, from which the iterative procedure refines the first stage decision variables.

C. Convergence

Figure 3 presents the convergence evolution of the full-EU instance, with Panels 3a and 3b reporting convergence over iterations and elapsed time, respectively. The regional initialization described in the previous subsection provides a tight initial upper bound, as evidenced by the best upper bound already being close to its final value at the first iteration. The lower bound rises rapidly during the first few iterations, and the optimality gap drops below 4% by iteration 25. After approximately 25 iterations the upper bound has largely stabilized and subsequent iterations are mainly devoted to improving the lower bound. The algorithm terminates after 49 iterations with an optimality gap of approximately 3%.

Panel 3b presents the same quantities over elapsed time. The total elapsed time is approximately 58 hours, with the upper bound stabilizing after roughly 27 hours. The remaining elapsed time is spent closing the gap by tightening the lower bound.

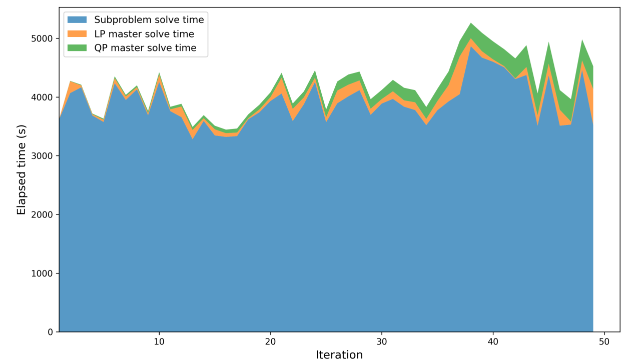

Figure 4. Time breakdown per iteration for the full-EU instance. The elapsed time per iteration is divided into three components: time to solve the subproblems, time to solve the level set master problem (QP), and time to solve the lower bound master problem (LP).

D. Computational run time

Figure 4 presents the elapsed time per iteration, broken down into three components. Over 49 iterations, the average time per iteration is approximately 70 minutes. The subproblem solve time dominates throughout the runs. As the iteration count increases, the time required to solve the level set master problem (QP) and the lower bound master problem (LP) grows, owing to the accumulation of cuts. Cut consolidation and cut sparsification mitigate this growth, keeping the subproblems as the computational bottleneck at every iteration. In this synchronous implementation, each iteration can proceed to the next master problem solve only after all workers have completed their assigned subproblems. As shown in the electronic supplement [29], the relative synchronization overhead ranges between 75% and 83% during the subproblem-solving phase[9], indicating substantial heterogeneity in worker completion times. This motivates asynchronous parallelization and improved load balancing strategies for allocating subproblems to workers.

[9]For each iteration, the relative synchronization overhead is computed as the difference between the slowest and fastest worker completion times, divided by the slowest worker completion time, during the subproblem-solving phase.

This article has been accepted for publication in IEEE Transactions on Power Systems. This is the author's version which has not been fully edited and content may change prior to final publication. Citation information: DOI 10.1109/TPWRS.2026.3725440

E. Comparison to the wait-and-see solution

As discussed in the initialization strategy of Section V-B, the monolithic wait-and-see problem for the full-EU instance is of excessive scale. Given our limited computational budget on ARIS, we therefore restrict this comparison to the Central-Western European (CWE) subsystem, for which the stochastic solution can be obtained via the proposed decomposition scheme and the wait-and-see solution can be computed by solving each scenario problem independently. Additional CWE results, including detailed breakdown tables, sensitivity analyses, and maps of invested capacities, are reported in the electronic supplement [29].

We quantify the cost of uncertainty through the so-called expected value of perfect information (EVPI) [9], defined as the additional cost incurred by the stochastic solution relative to the wait-and-see solution. The latter adapts first stage decisions to the realization of uncertainty, so a small EVPI suggests a limited role for uncertainty in investments. Based on the best incumbent solution of the stochastic problem, the estimated EVPI is approximately 3.9%, meaning that system costs would increase by roughly this amount if decisions were made without perfect foresight, highlighting the value of incorporating uncertainty into the optimization model. This additional cost stems almost entirely from higher annualized investment expenditure (6.44 billion € for the stochastic solution compared to 4.87 billion € for the wait-and-see), while economic dispatch costs remain comparable between the two solutions (27.18 and 27.49 billion €, respectively).

Figure 5 compares the stochastic and wait-and-see solutions for the CWE subsystem. Panel 5a presents the differences in the newly installed capacity between the wait-and-see solution and the stochastic solution. Under the stochastic solution, approximately 76.5 GW of new capacity is installed in total, compared to roughly 59.3 GW under the wait-and-see solution. Both wind and solar capacities are substantially higher in the stochastic solution, reflecting the additional hedging required when future conditions are unknown. By contrast, the wait-and-see solution is less aggressive in terms of investment, as it is able to perfectly anticipate the future conditions of the system and adapt its investment to these conditions.

The impact of these additional investments on system operation is illustrated in Panel 5b, which compares the renewable energy production of both solutions. The stochastic solution yields higher total renewable energy volumes, accompanied by higher production shedding, as the extra installed capacity generates output that cannot always be absorbed by the system. Nevertheless, this surplus ensures that sufficient renewable resources are available across a range of climate outcomes.

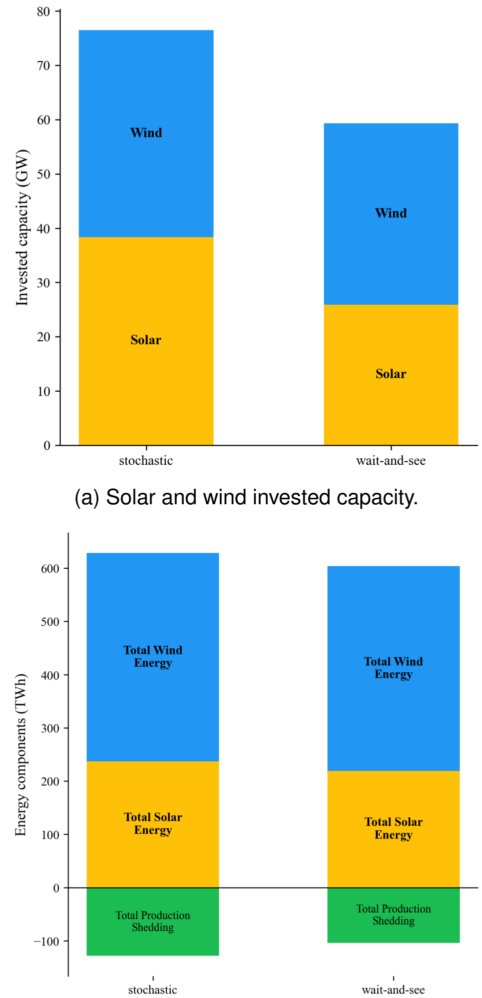

(a) Solar and wind invested capacity.

(b) Renewable energy production and shedding.

Figure 5. Comparison of the stochastic and wait-and-see solutions for the CWE subsystem. The upper panel shows the installed solar and wind capacity by technology, while the lower panel shows renewable energy production and shedding by technology.

VI. Conclusions

Renewable energy auctions can be an important policy instrument for steering investment in appropriate locations in the face of the ambitious renewable energy integration targets set by the European Union. Geographical coordination is essential, while network constraints, flexibility and storage, and uncertainty must all be accounted for. The formulation of this problem leads to a massive continuous stochastic capacity expansion problem, further complicated by a scenario-coupling renewable target band. We present a decomposition scheme specifically designed to exploit high performance computing infrastructure for tackling this problem on the scale of the full European power system. The algorithm relies on budget variable reformulations for the scenario-coupling target constraint and stored energy boundary variables for intertemporal storage coupling, dynamic level set stabilization to accelerate convergence, valid inequalities to strengthen the master problems, cut consolidation and cut sparsification to control their growth, and an initialization strategy that supplies a good starting point to the algorithm. The computational study on the full-EU instance demonstrates that the resulting problem can be solved to a 3% optimality gap within approximately 58 hours on 108 cores.

Future work will compare coordinated pan-European auction designs with country-level designs. Another line of research is to develop asynchronous parallelization strategies to mitigate synchronization delays caused by heterogeneous subproblem solution times and thereby reduce overall computation times. Finally, it is of interest to explore representative hour selection, grid aggregation, and scenario reduction techniques as means of reducing the dimensionality of the problem.

Acknowledgments

This work has received funding from the European Research Council (ERC) under the European Union's Horizon 2020 Research and Innovation Programme (Grant 850540). Moreover, this work was supported by computational time granted by the National Infrastructures for Research and Technology S.A. (GRNET S.A.) on the national HPC facility ARIS under project ID pr019005_fat. The authors would also like to thank Marilena Zampara (NTUA) for providing data for this work and Dr. Dimitrios Dellis (GRNET) for his continued support and assistance in facilitating access to GRNET computational resources. Finally, the authors thank the Editor of IEEE Transactions on Power Systems and the anonymous reviewers for their constructive feedback throughout the revision process.

References

[1] EC, "Europe's 2040 climate target and path to climate neutrality by 2050 building a sustainable, just and prosperous society," 2024. [Online]. Available: https://www.europarl.europa.eu/RegData/etudes/BRIE/2024/762862/EPRS_BRI(2024)762862_EN.pdf

[2] G. Thomassen, A. Fuhrmanek, R. Cadenovic, C. D. Pozo, and S. Vitiello, "Redispatch and congestion management," Joint Research Centre of the European Commission, Luxembourg (Luxembourg), Scientific analysis or review,Policy assessment,Anticipation and foresight KJ-NA-31-924-EN-N (online),KJ-NB-31-924-EN-N, 2024.

[3] M. Bjørndal and K. Jørnsten, "Zonal pricing in a deregulated electricity market," *The Energy Journal*, vol. 22, no. 1, pp. 51–73, 2001.

[4] Q. Lété, Y. Smeers, and A. Papavasiliou, "An analysis of zonal electricity pricing from a long-term perspective," *Energy Economics*, vol. 107, p. 105853, 2022.

[5] Q. Lété, Y. Smeers, and A. Papavasiliou, "Power generation investment under zonal electricity pricing with market-based re-dispatch," *The Energy Journal*, vol. 47, no. 1, pp. 131–164, 2026.

[6] I. Aravena, Q. Lété, A. Papavasiliou, and Y. Smeers, "Transmission capacity allocation in zonal electricity markets," *Operations Research*, vol. 69, no. 4, pp. 1240–1255, 2021.

[7] G. Thomassen and A. Fuhrmanek, "Locational price signals in europe," Joint Research Centre of the European Commission, Luxembourg (Luxembourg), Tech. Rep. KJ-01-25-525-EN-N, 2025.

[8] G. Thomaßen and A. Fuhrmanek, "Where to build renewables in europe? the benefits of locational auction design," *Energy Economics*, vol. 147, p. 108579, 2025.

[9] J. R. Birge and F. Louveaux, *Introduction to Stochastic Programming*, ser. Springer Series in Operations Research and Financial Engineering. New York, NY: Springer New York, 2011.

[10] J. A. Bloom, "Solving an electricity generating capacity expansion planning problem by generalized Benders' decomposition," *Operations Research*, vol. 31, no. 1, pp. 84–100, 1983.

[11] B. Gorenstin, N. Campodonico, J. P. da Costa, and M. Pereira, "Power system expansion planning under uncertainty," *IEEE Transactions on Power Systems*, vol. 8, pp. 129–136, 1993.

[12] 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.

[13] K. J. Singh, A. B. Philpott, and R. K. Wood, "Dantzig-Wolfe decomposition for solving multistage stochastic capacity-planning problems," *Operations Research*, vol. 57, no. 5, pp. 1271–1286, 2009.

[14] 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.

[15] J. Aghaei, N. Amjady, 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.

[16] 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, 2015.

[17] 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.

[18] 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.

[19] H. Teichgraeber, 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.

[20] M. Zampara, A. Visas, D. Ávila, L. Kliafas, and A. Papavasiliou, "A comparison of alternative algorithms for resolving stochastic expansion planning problems with endogenous adequacy targets," in *21st International Conference on the European Energy Market (EEM)*. IEEE, 2025, pp. 1–7.

[21] D. Ávila, A. Papavasiliou, M. Junca, and L. Exizidis, "Applying high-performance computing to the european resource adequacy assessment," *IEEE Transactions on Power Systems*, vol. 39, no. 2, pp. 3785–3797, 2024.

[22] T. Valencia Zuluaga, A. Musselman, J.-P. Watson, and S. S. Oren, "Parallel computing for power system climate resiliency: Solving a large-scale stochastic capacity expansion problem with mpi-sppy," *Electric Power Systems Research*, vol. 235, p. 110720, 2024.

[23] F. Pecci and J. D. Jenkins, "Regularized benders decomposition for high performance capacity expansion models," *IEEE Transactions on Power Systems*, vol. 40, no. 4, pp. 3105–3116, July 2025.

[24] F. Parolin, Y. Weng, P. Colbertaldo, and R. Macdonald, "Sectoral and spatial decomposition methods for multi-sector capacity expansion models," *Energy Conversion and Management*, vol. 358, p. 121356, 2026.

[25] A. Jacobson, F. Pecci, N. Sepulveda, Q. Xu, and J. D. Jenkins, "A computationally efficient Benders decomposition for energy systems planning problems with detailed operations and time-coupling constraints," *INFORMS Journal on Optimization*, vol. 6, no. 1, pp. 32–45, 2024.

[26] C. Wolf and A. Koberstein, "Dynamic sequencing and cut consolidation for the parallel hybrid-cut nested l-shaped method," *European Journal of Operational Research*, vol. 230, no. 1, pp. 143–156, 2013.

[27] G. Thomaßen, "Introducing locational criteria to renewable auctions," *IEEE Transactions on Energy Markets, Policy and Regulation*, pp. 1–10, 2025.

[28] A. Papavasiliou, *Optimization Models in Electricity Markets*. Cambridge University Press, 2024.

[29] A. Visas, D. Ávila, A. Papavasiliou, G. Thomassen, and D. Pozo, "Electronic supplement: Decomposition techniques for model-based bid selection in renewable auctions," 2026. [Online]. Available: https://ap-rg.eu/wp-content/uploads/2026/08/Electronic_Supplement__Decomposition_techniques_for_model_based_bid_selection_in_renewable_auctions_final.pdf

[30] R. Rahmaniani, T. G. Crainic, M. Gendreau, and W. Rei, "The benders decomposition algorithm: A literature review," *European Journal of Operational Research*, vol. 259, no. 3, pp. 801–817, 2017.

[31] V. Zverovich, C. I. Fábián, E. F. D. Ellison, and G. Mitra, "A computational study of a solver system for processing two-stage stochastic LPs with enhanced Benders decomposition," *Mathematical Programming Computation*, vol. 4, no. 3, pp. 211–238, 2012.

[32] L. Göke, F. Schmidt, and M. Kendziorski, "Stabilized Benders decomposition for energy planning under climate uncertainty," *European Journal of Operational Research*, vol. 316, no. 1, pp. 183–199, 2024.

[33] J. Linderoth and S. Wright, "Decomposition algorithms for stochastic programming on a computational grid," *Computational Optimization and Applications*, vol. 24, pp. 207–250, 2003.

[34] C. Lemaréchal, A. Nemirovskii, and Y. Nesterov, "New variants of bundle methods," *Mathematical Programming*, vol. 69, no. 1, pp. 111–147, 1995.

[35] N. Stevens and A. Papavasiliou, "Application of the level method for computing locational convex hull prices," *IEEE Transactions on Power Systems*, vol. 37, no. 5, pp. 3958–3968, 2022.

[36] J. Kronqvist, D. E. Bernal, and I. E. Grossmann, "Using regularization and second order information in outer approximation for convex MINLP," *Mathematical Programming*, vol. 180, pp. 285–310, 2020.

[37] H. Zhang, I. E. Grossmann, K. McKinnon, B. R. Knudsen, R. G. Nava, and A. Tomasgard, "Integrated investment, retrofit and abandonment energy system planning with multi-timescale uncertainty using stabilised adaptive benders decomposition," *European Journal of Operational Research*, vol. 325, no. 2, pp. 261–280, 2025.

[38] S. Trukhanov, L. Ntaimo, and A. Schaefer, "Adaptive multicut aggregation for two-stage stochastic linear programs with recourse," *European Journal of Operational Research*, vol. 206, no. 2, pp. 395–406, 2010.

[39] R. Rahmaniani, T. G. Crainic, M. Gendreau, and W. Rei, "An asynchronous parallel Benders decomposition method for stochastic network design problems," *Computers & Operations Research*, vol. 162, p. 106459, 2024.

[40] G. K. Saharidis, M. Boile, and S. Theofanis, "Initialization of the Benders master problem using valid inequalities applied to fixed-charge network problems," *Expert Systems with Applications*, vol. 38, no. 6, pp. 6627–6636, 2011.

[41] T. Brown, J. Hörsch, and D. Schlachtberger, "PyPSA: Python for Power System Analysis," *Journal of Open Research Software*, vol. 6, no. 1, p. 4, Jan. 2018.

[42] ENTSO-e, "ERAA 2021 Modelling data." [Online]. Available: https://www.entsoe.eu/eraa/2021/modelling-data/

Alexandros Visas (Student Member, IEEE) received the M.Eng. degree in Electrical and Computer Engineering from the University of Patras and the M.Sc. degree in Energy Production and Management from the National Technical University of Athens (NTUA). He has worked at the Greek transmission system operator, the Independent Power Transmission Operator (IPTO or ADMIE), focusing on power system planning, renewable energy integration, and battery storage integration. He is currently a Ph.D. candidate with the School of Electrical and Computer Engineering, NTUA, under the supervision of Prof. Anthony Papavasiliou. His research interests include power system planning, renewable energy and storage, and stochastic optimization.

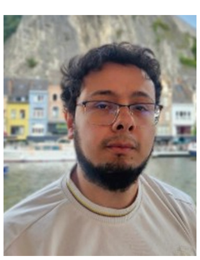

Daniel Ávila received the bachelor's and master's degrees in mathematics from Universidad de los Andes and the Ph.D. degree from the Université catholique de Louvain, under the supervision of Prof. Anthony Papavasiliou. He later continued at the Université catholique de Louvain as a Postdoctoral Researcher with the Center for Operations Research and Econometrics (CORE). He is currently an Energy Consultant with Arktos Analytics and a Research Associate with the National Technical University of Athens (NTUA), within the research group of Prof. Anthony Papavasiliou. His research interests include convex optimization, stochastic programming, and renewable energy integration in power systems.

Anthony Papavasiliou (Senior Member, IEEE) is an associate professor at the department of Electrical and Computer Engineering at the National Technical University of Athens. Formerly, he was an associate professor at the Center for Operations Research and Econometrics (CORE) at the Université catholique de Louvain, where he held the ENGIE Chair and a Francqui Research Professorship. He completed his PhD and post-doctoral research at the department of Industrial Engineering and Operations Research at UC Berkeley and his undergraduate studies at the National Technical University of Athens, Greece. In 2026 Anthony received the Dieter Schwarz Courageous Research Grant of the Institute of Advanced Studies at the Technical University of Munich. In 2021 he received the Bodossaki Foundation Distinguished Young Scientist award in Applied Sciences. He is a recipient of an ERC Starting Grant in 2019. He is serving as an associate editor for Operations Research and has served as an Associate Editor for the IEEE Transactions on Power Systems, as well as in the technical program committees of various international conferences. He has received the 2015 INFORMS best publication award in energy. He has served as a consultant for transmission system operators, regulatory authorities, power exchanges, and utilities on a number of topics related to electricity markets and power system operations.

Georg Thomassen holds an M.Sc. degree from the Technical University of Berlin and a Ph.D. degree from Leipzig University, where his research focused on electricity market design for the energy transition. He is currently Technical Lead for Electricity Market and Infrastructure Modeling at the European Commission's Joint Research Centre (JRC). Previously, he worked as an Analyst at Agora Energiewende. His research interests include electricity market design, resource adequacy, and the integration of renewable energy and flexibility into power systems.

David Pozo (Senior Member, IEEE) is a Scientific Project Officer at the Joint Research Centre (JRC) of the European Commission, where his activities combine energy-system analysis, mathematical optimization, and policy support. Before joining the JRC in 2022, he was an Assistant Professor at the Skolkovo Institute of Science and Technology (Skoltech). He held postdoctoral fellowships at the Pontifical Catholic University of Rio de Janeiro, Brazil, and the Pontifical Catholic University of Chile, Santiago, Chile. He received the Ph.D. degree in electrical engineering from the University of Castilla-La Mancha, Spain. His research spans operations research, economics, and power systems, with applications to electricity markets, power-system planning and operation, decision-making under uncertainty, flexibility, and low-carbon energy systems. His distinctions include the 2024 UCLM Alumni Research Award, the 2023 JRC Award for Excellence in Science for Policy, the 2023 Xpress Best Paper Award, the 2021 IEEE TPWRS Best Paper Award, and the 2020 INFORMS Best Publication Award in Energy. He serves as an Associate Editor of the IEEE Transactions on Power Systems and is also a member of the technical program committees of several international conferences.