Available online at www.sciencedirect.com

ScienceDirect

IFAC Papers Online CONFERENCE PAPER ARCHIVE

IFAC PapersOnLine 59-9 (2025) 1–6

Distributed Virtual Power Plant Optimization via Scenario Approach ⋆

Alessandro Del Duca * **Fredy Ruiz** *

* *Dipartimento di Elettronica, Informazione e Bioingegneria, Politecnico di Milano, 20133, Milano, Italy. (e-mail: alessandro.delduca@polimi.it, fredy.ruiz@polimi.it)*

Abstract: The shift towards decentralised energy systems has given rise to Virtual Power Plants (VPPs), essential for managing Distributed Energy Resources. This study proposed a distributed management architecture for VPPs that resiliently provides aggregated power production to the main grid, irrespective of uncertainties in non-dispatchable resources. Exploiting a scenario optimisation formulation and a Lagrangian decomposition of the problem, we develop a distributed stochastic optimisation framework for VPP management that scales with the number of agents. The method's performance is tested across various network sizes and real generation and consumption data, showing that the constraint violation probability on the minimum service level remains constant as the number of agents increases, while maintaining a fixed amount of considered scenarios.

Keywords: Energy Management, Virtual Power Plants, Smart Grids, Distributed Optimisation, Multi-agent systems, Randomised methods.

1. INTRODUCTION

The transition towards decentralised energy systems has led to the emergence of Virtual Power Plants (VPPs) as a key solution for managing Distributed Energy Resources (DERs), such as Photovoltaic (PV) plants, Battery Energy Storage Systems (BESSs), Electric Vehicles (EVs) and Flexible Loads (FL), as a single controllable power production unit. Unlike traditional power plants, VPPs do not rely on physical infrastructures to offer services to the power grid but leverage advanced control algorithms and communication networks to optimise energy dispatch of distributed assets. Consequently, a VPP can participate in the primary and ancillary energy market as a generation plant, enabling aggregated DERS to provide generation and ancillary services.

A common approach to optimising VPP operations involves a centralised architecture. In a centralised VPP, a central unit makes all decisions about the operating conditions of each DER. Yang et al. (2020) present a centralised VPP capable of aggregating EVs to participate in the short-term energy market. Diaz et al. (2018) propose a centralised strategy to manage EV charging stations based on a Model Predictive Control (MPC) strategy. Hannan et al. (2019) describes a centralised schedule controller for a VPP based on Particle Swarm Optimisation. The drawback of centralised architectures is that the DERs must communicate with the central operator, possibly exposing private information about their parameters and operations. Moreover, as the number of DERs increases, the communication requirements for effective coordination can become quite complex.

Another critical aspect in the optimal management of DERs is addressing the uncertainty in power production and consumption and the Transmission System Operator (TSO) service requests. The literature highlights two major families of methods to handle uncertainty: two-stage stochastic optimisation and Scenario optimisation. Liang et al. (2019) propose a VPP considering correlated demand response (CDR), formulating the problem as a risk-constrained optimisation problem solved with a two-stage stochastic program. Similarly, Vahedipour-Dahraie et al. (2021) propose optimal energy and reserves scheduling for VPPs based on a two-stage risk-constrained stochastic program. Baringo et al. (2019) present a VPP architecture with uncertainty in wind production, prices and TSO requests. The VPP solves an adjustable robust optimisation problem by exploiting the concept of uncertainty budget. Finally, Naughton et al. (2021) present a two-level hierarchical VPP architecture based on the scenario theory form Campi and Garatti (2011); the model considers uncertainty in PV production, power consumption and prices and runs at two different time scales, to compensate for the uncertainty.

In this work, we present an architecture for a VPP based on the distributed stochastic optimisation framework proposed in Del Duca et al. (2025), applied to a network of interconnected microgrids. Our proposed VPP model aggregates different microgrids, which are technically capable of working as a small-scale power grid, as in Del Duca et al. (2024), but may fail to participate in the primary market. The VPP enable them to participate in the primary market as a single cohesive generation unit. In our ap-

⋆ This research has been supported by the Italian Ministry of University and Research (MIUR) under the PNRR program and by Hitachi Energy Italy S.p.A co-financer of the grant.
This research has been supported by the Italian Ministry of University and Research under grant "Learning-based Model Predictive Control by Exploration and Exploitation in Uncertain Environments" (PRIN PNRR 2022 fund, ID P2022EXP2W).

10.1016/j.ifacol.2025.08.103

proach, the scenario theory is tailored to solve distributed optimisation problems with uncertain global constraints. The VPP optimises the microgrid operation through the Augmented Lagrangian distributed algorithm from Falsone and Prandini (2023), ensuring the local balance of all microgrids and a minimum level of aggregated power capacity to comply with requests from the TSO. This architecture offers several advantages. Microgrids maintain their independence as legal entities with decision-making capabilities and do not have to disclose private information about their assets, such as operational constraints and costs. Instead, they only need to communicate the dual variables of the optimisation problem, as explained in section 4. The lack of a single central unit decreases the communication burden of the VPP operation. The VPP becomes naturally robust to microgrid disconnections, communication issues, and network reconfiguration. We test the VPP with extensive simulations in section 5, verifying the framework characteristics against networks with different sizes, confirming that it is possible to decouple the uncertainty handling from the network complexity, supporting the claims previously made in Del Duca et al. (2025).

In section 2, we define the VPP optimization problem; in section 3 we provide probabilistic guarantees on constraints violation levels; in section 4 we present an algorithm to solve the VPP problem in a distributed way; in section 5 we show simulation results; in section 6 we make some final remarks and future work.

2. VIRTUAL POWER PLANT PROBLEM DEFINITION

2.1 Virtual Power Plant modelling

We consider a network of m microgrids, also called agents, interconnected through a graph $\mathcal{G} = (\mathcal{V}, \mathcal{E})$ with $\mathcal{V} = \{1..m\}$ being the set of vertices (microgrids) and $\mathcal{E} = \{(i,j) : i \in \mathcal{N}_j \quad \& \quad j \in \mathcal{N}_i\}$ being the set of edges. The graph's edges $(i,j) \in \mathcal{E}$ serve two distinct functions. Microgrids connected by an edge may exchange power $P_{(i,j)}$ through a physical bus; the $i-th$ microgrid can also exchange information about its dual variables through a communication link to its set of neighbours $\mathcal{N}_i$.

2.2 Microgrid modeling

A microgrid is a decision-making entity (agent) that acts on local private decision variables $x_i(t)$ representing the power dispatch of the microgrid's assets during the optimisation horizon $t \in [0, T-1]$. Each microgrid is defined as a discrete-time LTI system with sampling time $\Delta t = 1h$; it is formed by a dispatchable generation unit, a non-dispatchable solar generation unit, and a load that consumes power. Some microgrids may have a Battery Energy Storage System (BESS) to improve flexibility and provide energy storage capabilities.

Dispatchable generation units: A dispatchable generation unit can generate a controllable amount of power $P^i_{gen} \in [0, \overline{P}^i_{gen}]$, with $\overline{P}^i_{gen}$ its maximum capacity. The production cost $c^i_{gen}(t)$ increases quadratically with $P^i_{gen}(t)$:

$$c^i_{gen} = \sum_{t=1}^{T} c^i_{gen}(t) = \sum_{t=1}^{T} (\alpha_i \cdot (P^i_{gen}(t))^2 + \beta_i \cdot P^i_{gen}(t) + \gamma_i)$$

The generation units are also subject to maximum ramp-rate constraints:

$$|P^i_{gen}(t) - P^i_{gen}(t-1)| \leq \overline{P}^i_{ramp}$$

Battery Energy Storage Systems: The State of Charge of a BESS, $s^i(t)$, represents the level of energy stored in it; the BESS's dynamics are described as an LTI model with the power exchange with the microgrid $P^i_{bess}(t)$ as input, that is:

$$\begin{cases} s^i(t+1) = s^i(t) - \dfrac{\Delta t}{Q_i} \cdot P^i_{bess}(t) \\ s^i(0) = s^i(0) \end{cases}$$

where s^i_0) is the known BESS's initial State of Charge (SoC). $P^i_{bess}(t)$ is assumed to be positive when the battery is discharged and negative when charged. In addition, $P^i_{bess}(t)$ must lay inside safe operational constraints $P^i_{bess}(t) \in [\underline{P}^i_{bess}, \overline{P}^i_{bess}]$.

The SoC cannot exceed operational limits $s^i(t) \in [\underline{s}; \bar{s}]$, while the initial and final SoC must be equal to ensure enough capacity for the next day: $s^i(0) = s^i(T-1)$. Each BESS is associated with a cost to prevent excessive degradation due to fast charge and discharge operations:

$$c^i_{bess} = \sum_{t=0}^{T-1} (\delta_i \cdot (P^i_{bess}(t) - P^i_{bess}(t-1))^2)$$

Power exchange with neighbours microgrids: The i-th microgrid can exchange power with all its neighbours $j \in \mathcal{N}_i$ through the bus $(i,j) \in \mathcal{E}$. The power $P_{(i,j)}(t)$ must lay inside operational constraints: $P_{(i,j)}(t) \in [\underline{P}_{(i,t)}, \overline{P}_{(i,t)}]$; the cost of these exchanges depend on the power losses, and therefore it scales quadratically with $P_{(i,j)}(t)$ thought a constant factor $R_{(i,j)}$:

$$c_{(i,j)} = \sum_{t=0}^{T} (R_{(i,j)} \cdot P^2_{(i,j)}(t))$$

With this penalisation term, we can encode a linearised power flow approximation into the optimisation, as described in Naughton et al. (2021).

Non-dispatchable production and consumption: Each microgrid has local solar production $\hat{P}^i_{pv}(t)$ and local consumption $\hat{P}^i_l(t)$, for which forecasts are available.

In brief, the microgrid has to satisfy a local power balance of the form:

$$P^i_{gen}(t) + P^i_{bess}(t) + \hat{P}^i_{pv}(t) - \hat{P}^i_l(t) + \sum_{j \in \mathcal{N}_i} P^i_j(t) \geq 0 \quad (1)$$

In eq. (1), the sign is positive if the power is produced by the dispatchable and PV generators, made available by discharging the BESS or imported from neighbours; the sign is negative if the power is consumed by the load, used to charge the BESS, or exported to other neighbours. This constraint must be equal to zero if the TSO does not make an aggregated request. Instead, it is greater than zero for some microgrids when the TSO asks for a ramp-up since the non-zero contributions will aggregate to the total VPP power export.

Based on the previous description of the assets, the decision variables of each microgrid are

$$x_i(t) = [P^i_{gen}(t), P^i_{bess}(t), s^i(t), \{P^i_j(t)\}_{j\in\mathcal{N}_i}] \ \forall t,$$

Note that all the operational constraints of the microgrid can be rewritten compactly as a convex set: $x_i \in X_i$. Each microgrid should optimise operation costs:

$$f_i(x_i) = c^i_{gen} + c^i_{bess} + 0.5 \cdot \sum_{j\in\mathcal{N}_i} c_{(i,j)}$$

while remaining within operational ranges independently of other microgrids. Therefore, if the microgrids were not interconnected, each agent would have determined the optimal power schedule for their assets x^*_i as follows:

$$\mathcal{P}_i : \ min_{x_i} f_i(x_i) \ \ s.t. \quad x_i \in X_i$$

2.3 Coupling constraints

The following constraints encode the coupling between microgrids.

Power-flow coupling constraints: Each microgrid is an independent entity with its own set of decision variables $P^i_j(t) \ \forall j \in \mathcal{N}_i$ representing the power exchange with neighbour microgrids. To ensure power balance across the network, for each edge $(i,j) \in \mathcal{E}$, the following equality constraint must hold:

$$P_{(i,j)} = P^i_j(t) = -P^j_i(t) \ \forall (i,j) \in \mathcal{E} \ \forall t \tag{2}$$

with the positive sign indicating that the power is flowing from j to i. This constraint enforces that if j imports power $P^j_i(t)$ from i, i exports the same amount of power $P^i_j(t)$ to j at every time-step t. With some algebraic manipulation, we can express these constraints as a linear combination of all the microgrids' variables $x_i(t)$:

$$\sum_{i=1}^{m} \hat{A}_i \cdot x_i(t) = \underline{0}$$

The $|\mathcal{E}| \times |\mathcal{N}_i|$ matrix A_i assigns the decision variable $P^i_j(t)$ to the correct constraint $(i,j) \in \mathcal{E}$. Therefore, each row of the constraint represents the balance for a single link, and only the variables of the $i-th$ and $j-th$ microgrid are involved. This constraint is then extended over the optimisation horizon as

$$\sum_{i=1}^{m}((\hat{A}_i \otimes I_T) \cdot [x^T_i(0)..x^T_i(T)]^T = \sum_{i=1}^{m} A_i \cdot x_i = \underline{0}$$

where I_T is the identity matrix with dimension $T \times T$, $\otimes$ is the Kronecker product and $(\hat{A}_i \otimes I_T) = A_i$.

Uncertain constraint on aggregated power generation: The Virtual Power Plant must always provide enough power as a generation unit; however, due to the uncertain nature of the aggregated PV production $\hat{P}_{pv}(t)$, and load consumption $\hat{P}_l(t)$:

$$\hat{P}_{pv}(t) = \sum_{i=1}^{m} P^i_{pv}(t), \quad \hat{P}_l(t) = \sum_{i=1}^{m} P^i_l(t)$$

The amount of power effectively produced is uncertain. It depends on the realization of the random variable $\delta(t) = \hat{P}_{pv}(t) - \hat{P}_l(t) \in \Delta$, where Δ is the (unknown) set of all possible realizations of the aggregated net-demand. The amount of power the VPP can deliver to the TSO is:

$$P_n(t,\delta) = \sum_{i=1}^{m}(P^i_{gen}(t) + P^i_{bess}(t)) + \delta(t).$$

The latter constraint arises from the power balance across the network, notably, according to Eq. (2), the total power exchanges among the agents sum to zero, implying that we do not need to account for internal exchanges when calculating the power the VPP delivers to the primary Grid.

The Transmission System Operator dictates this net production as an exogenous signal $P^{ref}_n(t)$ that the VPP must comply with $P_n(t,\delta) \geq P^{ref}_n(t) \forall t$.

In practice, it is not convenient to robustify the solution against the entirety of the uncertain set Δ because it will result in excessive conservativeness. To alleviate this problem, we impose the following chance-constraint:

$$\mathbf{P}_\delta\{\delta \in \Delta : P_n(t,\delta) - P^{ref}_n(t) \geq 0\} \geq 1-\epsilon \tag{3}$$

Eq. (3) imposes that the probability of satisfying the TSO request, that is, providing an aggregated power $P_n(t.\delta)$ greater than $P^{ref}_n(t)$ must be greater than $1-\epsilon$, where $\epsilon \in (0,1)$ is called the risk parameter. Similarly to eq. (2) we can rewrite this constraint for $t \in [0, T-1]$ compactly:

$$\mathbf{P}\{\delta \in \Delta : \sum_{i=1}^{m} H_i \cdot x_i + H_\delta \delta \geq 0\} \geq 1-\epsilon \tag{4}$$

for proper matrices H_i and H_δ.

2.4 Centralized optimization problem

In this work, the Virtual Power Plant optimisation is performed over a full day (day-ahead) $T = 24h$ with sampling time $\Delta t = 1h$. However, the proposed approach remains valid for different horizons and sampling periods.

The VPP executes the optimisation each instant in a receding-horizon fashion. At each time step, the centralised problem solved is:

$$\begin{aligned} \mathcal{P} : & \ min_{\{x_i\}_{i=1}^m} \sum_{i=1}^{m} f_i(x_i) \ \ s.t. \\ & x_i \in X_i \ \forall i = 1..m \\ & \sum_{i=1}^{m} A_i \cdot x_i = 0 \\ & \mathbf{P}\{\delta \in \Delta : \sum_{i=1}^{m} H_i \cdot x_i + H_\delta \cdot \delta \geq 0\} \geq 1-\epsilon \end{aligned} \tag{5}$$

3. UNCERTAINTY HANDLING AND PROBABILISTIC GUARANTEES

The stochastic optimisation problem in Eq. (5) is not easy to solve, not even in a centralised way, due to the chance-constraint Eq. (3). To handle such constraints, in this work we rely on the framework presented in Del Duca et al. (2025) and approximate the uncertain chance-constraint with a finite set of sampled constraints, where the uncertainty δ is sampled from historical data, and each sample δ_k is called a scenario. The stochastic constraint (4) is transformed into a set of N linear constraints as

$$\sum_{i=1}^{m} H_i \cdot x_i + H_\delta \cdot \delta_k \geq 0 \ \ \forall k = 1, \ldots, N \tag{6}$$

or equivalently

$$\bigcap_{k=1}^{N}(\mathcal{F}_k) \; with \; \mathcal{F}_k := \{x : \sum_{i=1}^{m} H_i \cdot x_i + H_\delta \cdot \delta_k \geq 0\} \quad (7)$$

By choosing the number of scenarios N high enough, the solution of the approximated optimisation program $\{x_i^*\}_{i=1}^m$ has a-priori probabilistic guarantees under unseen scenarios δ_{N+1}. Given the structure of constraint (4), based on Del Duca et al. (2025), the following theorem holds:

Theorem 1. Given a risk parameter $\epsilon \in (0,1)$ and a confidence parameter $\beta \in (0,1)$, the solution $x^* = [x_1^T, \dots x_m^T]^T$ to the scenario optimization problem (5), replacing constraint (4) by (6), is so that:

$$\mathbf{P}[\delta \in \Delta : \sum_{i=1}^{m} H_i \cdot x_i^* + H_\delta \cdot \delta < 0] \leq \epsilon \quad (8)$$

with high confidence $1-\beta$, if N is such that the relation

$$\epsilon = 1 - \sqrt[N-\rho]{\frac{\beta}{\binom{N}{\rho}}} \quad (9)$$

is satisfied, and $\rho = T$ is the Helly's dimension of the problem.

Given a dataset of scenarios $S^N = \{\delta_1..\delta_N\} \in \Delta^N$ independently chosen from the uncertain set Δ, Theorem 1 states that the probability of violation remains inside the prescribed risk level ϵ with high confidence $1-\beta$. The number of scenarios N needed to guarantee a prescribed violation level ϵ depends on Helly's dimension ρ of the problem, which in a VPP setting scales with the network size m; this scaling could make the scenario approximation hard to compute for large networks. However, we exploit the structure of the chance-constraint in Eq. (6), using the *Lemma 2* in Del Duca et al. (2025), to compute N using a much lower bound on ρ which is equal to the dimension r of the row-span of the constraint matrix H_i, equals the optimisation horizon $\rho \leq r = T$. This new bound separates the uncertainty handling from the network complexity.

4. DISTRIBUTED OPTIMIZATION

The scenario program resulting in substituting the chance-constraint in (5) with (6) becomes a quadratic program solvable with standard techniques in a centralised fashion. Nonetheless, solving the problem in such a way has the disadvantages stated in 1.

However, problem (5) can be decomposed into m smaller sub-problems, one for each microgrid. We can use a distributed scheme as in Falsone and Prandini (2023) to solve the now deterministic optimisation program. The algorithm is tailored to track equality and inequality coupling constraints of the same structure of the power-flow balance (2) and the scenario chance-constraint (6). The procedure is summarised in Algorithm 4.

The algorithm has an initialisation part:

- *Local scenario extraction (2):* Each microgrid extracts $N_i \geq \lceil \frac{N}{m} \rceil$ scenarios of PV generation and load for the complete horizon T from recorded data.
- *Scenario broadcast (3):* Each microgrid broadcasts its local scenarios to assemble the data-set $\tilde{D}$.
- *Initialization (4-6):* Each microgrid initializes its dual variables λ_i, μ_i, the initial primal solution $x_i \in X_i$ and the constraint residuals $g_i(x_i), \sigma_i, d_i$.

Algorithm 1 Augmented Lagrangian Tracking (ALT)

1: **Initialization:**
2: $\tilde{D}_i := \{\delta_1..\delta_{\lceil \frac{N}{m} \rceil + 1}\} \forall i = 1..m$ ▷ *Generate scenarios*
3: $\tilde{D} = \bigcap_{i=1}^m \tilde{D}_i$ ▷ *Broadcast scenarios*
4: $x_i^0 \in \mathbb{R}^{n_i}$, $\sigma_i^0 \in \mathbb{R}_+^q$ ▷ *Initialize variables*
5: $\lambda_i^0 \in \mathbb{R}^p$, $d_i^0 = -(A_i x_i^0 - b_i)$
6: $\mu_i^0 \in \mathbb{R}^q$, $g_i^0 = -(h_i(x_i^0) + \sigma_i^0)$
7: **while** not converged **do**
8: *Consensus step*
9: $\ell_i^\tau = \sum_{j \in \mathcal{N}_i} w_{ij} \lambda_j^\tau, \quad m_i^\tau = \sum_{j \in \mathcal{N}_i} w_{ij} \mu_j^\tau$
10: $\delta_i^\tau = \sum_{j \in \mathcal{N}_i} w_{ij} d_j^\tau, \quad \gamma_i^\tau = \sum_{j \in \mathcal{N}_i} w_{ij} g_j^\tau$
11: *Primal update*

$$x_i^{\tau+1} \in \arg\min_{x_i \in X_i} \Big\{ f_i(x_i) + (\ell_i^\tau)^\top A_i x_i + \frac{c}{2} \|A_i x_i - A_i x_i^\tau - \delta_i^\tau\|^2 + \frac{1}{2c} \|\max\{m_i^\tau + c(h_i(x_i) - h_i(x_i^\tau) - \sigma_i^\tau - \gamma_i^\tau), 0\}\|^2 \Big\}$$

12: *Dual update*
13: $\sigma_i^{\tau+1} = \gamma_i^\tau - h_i(x_i^{\tau+1}) + h_i(x_i^\tau) + \sigma_i^\tau - \frac{1}{c} m_i^\tau$
14: $\sigma_i^{\tau+1} = \max\{\sigma_i^{\tau+1}, 0\}$
15: $d_i^{\tau+1} = \delta_i^\tau - A_i x_i^{\tau+1} + A_i x_i^\tau$
16: $g_i^{\tau+1} = \gamma_i^\tau - (h_i(x_i^{\tau+1}) + \sigma_i^{\tau+1}) + (h_i(x_i^\tau) + \sigma_i^\tau)$
17: $\lambda_i^{\tau+1} = \ell_i^\tau - c d_i^{\tau+1}, \quad \mu_i^{\tau+1} = m_i^\tau - c g_i^{\tau+1}$
18: $\tau \leftarrow \tau + 1$s
19: **end while**

Then, the algorithm finds a solution for the primary variables by iteratively performing primal and dual steps. At each iteration, $\tau = 1, \dots, \tau_{max}$, in parallel:

- *Consensus step (8-10):* Each microgrid communicates its estimated dual variables λ_i, μ_i and tracking residuals $g_i(x_i), \sigma_i, d_i$ to its neighbours. Then, each microgrid computes the local average of such variables by weighting each neighbour variable according to the weights a_i^j, which reflect the network's topology.
- *Primal solution step (11):* Each microgrid computes the candidate primal solution, solving the local optimisation problem, augmented with the proximal costs necessary to handle the coupling constraints.
- *Sub-gradient ascent step (12-16):* Each microgrid computes the new estimate for the tracking residuals based on the latest candidate primal solution. It uses the new estimate to perform a gradient step on the estimate of the local dual variables l_i, m_i.

The following assumptions on the graph topology are necessary to guarantee that the distributed scheme can reach a consensus, see Falsone and Prandini (2023).

Hypothesis 2. (connectivity). Graph $(\mathcal{V}, \mathcal{E}_\infty)$ is strongly connected, i.e., for any pair of nodes, there exists a path of directed edges that connects them. Moreover, there exists a time-period $\delta t > 1$ such that for every $(i,j) \in \mathcal{E}_\infty$ agent i receives information from a neighbour $j \in \mathcal{N}_i$ at least once every δt iterations.

Hypothesis 3. (Information mixing). It exists $\eta \in (0,1)$ such that for all $i,j \in \{1,\dots,m\}$ and all $t \geq 0$,

$$a_{ij}^{(t)} \in [0,1), \quad a_{ii}^{(t)} \geq \eta, \quad \wedge \quad a_{ij}^{(t)} > 0 \Longrightarrow a_{ij}^{(t)} \geq \eta.$$

Moreover, for all $t \geq 0$,

(1) $\sum_{j=1}^{m} a_{ij}^{(t)} = 1$ for all $i = 1, \ldots, m$,
(2) $\sum_{i=1}^{m} a_{ij}^{(t)} = 1$ for all $j = 1, \ldots, m$.

The first assumption states that information can reach every node in the network; the second ensures that the information is mixed at a non-diminishing rate, meaning that all the agents equally influence their neighbours. These assumptions ensure the algorithm converges to the centralised solution, so the reasoning in section 3 is valid. The algorithm converges even under time-varying communication networks, making the VPP robust to microgrid temporary disconnections.

We also stress the importance of relying on a bound on Helly's dimension, independent of the number of agents in Eq. (9). If this were not the case, each agent would have to communicate a vector of dual variables λ, whose dimension would grow with the number of agents m, limiting its applicability on large-scale systems.

5. SIMULATION RESULTS

To verify the proposed VPP architecture, we test it on four different networks with increasing size: $(m = 8, |\mathcal{E}| = 16)$, $(m = 20, |\mathcal{E}| = 70)$, $(m = 50, |\mathcal{E}| = 125)$, $(m = 100, |\mathcal{E}| = 150)$ and a desired maximum violation parameter $\epsilon = 10\%$, that implies $N = 362$ scenarios in eq. (9). We repeated each network's daily simulation 100 times, choosing different scenario databases S^N and unseen scenarios to test the violation probability. The data is available in the open-source database in Vink et al. (2019).

After computing the optimal distributed solution x_N^* by solving the N-scenario program in eq. (5) with the distributed algorithm 4, we test the optimal scheduling against unseen scenarios; for our analysis, we chose to test each solution against $Mc = 1000$ unseen scenarios.

5.1 Aggregated power production probabilistic guarantees

In the simulations, the TSO does not request power from the VPP in the interval $t \in (0, 4)$, i.e. it is sufficient that the VPP does not import power from the main grid. Then, starting from $t = 4h$, the reference power production $P_n^{ref}(t)$ increases to a value of $P_n^{ref}(t) = 400KW$, and the VPP must adjust the aggregated production accordingly. The request is the dashed red line in Fig. 1.

Figure 1 shows a single instance of the aggregated power production $P_n(t, \delta)$, the mean value (blue line) with the $1 - \epsilon$ level (light blue area). The $1 - \epsilon$ level represents the area containing the $1-\epsilon$ mass of the scenarios, i.e. the area containing the scenarios that do not violate the constraint under the a-priori guaranteed level of risk. Instead, some examples of solutions violating constraints are depicted as red dotted lines. It can be noticed that the VPP can correctly schedule the aggregated power production while minimising the costs of the entire network of agents and ensuring the satisfaction of local operational constraints.

To verify that the a-priori violation levels are satisfied for all the networks, we define the violation probability as:

$$V_N(x^*) = \mathbf{P}[\delta \in \Delta : \sum_{i=1}^{m} H_i \cdot x_i^* + H_\delta \cdot \delta < 0] \qquad (10)$$

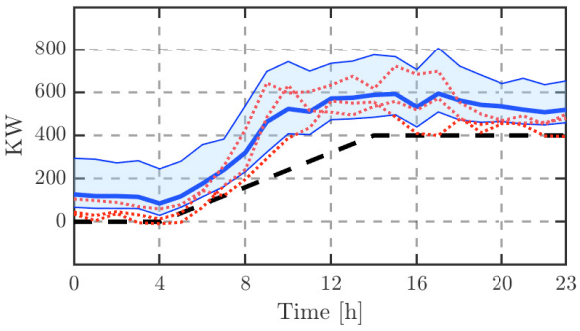

Fig. 1. Aggregated power production $P_n(t, \delta)$: mean value (blue line), 1-ϵ level (light blue area), violating unseen scenarios (dotted red line). Minimum aggregated power required $P_n^{ref}(t)$ (dashed black line)

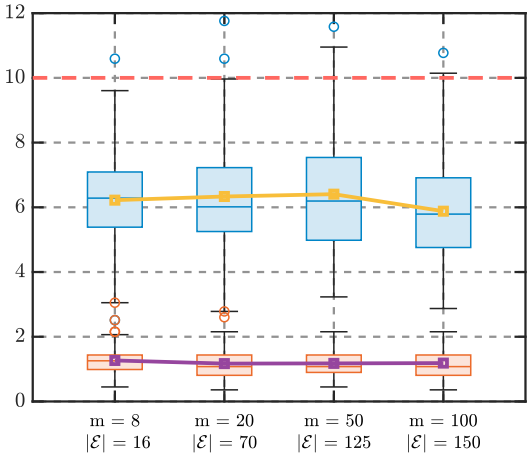

Fig. 2. Violation probability $V_N(x^*, \delta)$, in red $N = 362$, in blue $N = 100$, yellow and purple line: average value of the violation

The violation probability $V_N(x^*)$ can be approximated by sampling the violation levels for each daily simulation.

The resulting sampled violation $\hat{V}_N(x^*)$ is:

$$\hat{V}_N(x^*) = \frac{1}{Mc} \sum_{j=1}^{Mc} [\mathbf{1}_{V_N}(x^*, \delta_j)] \qquad (11)$$

where $1_{V_n}(x)$ is the indicator function, indicating if the scenario δ_i violates the constraint.

$$\mathbf{1}_{V_n}(x^*, \delta_j) := \begin{cases} 1 & \text{if } (\sum_{i=1}^{m} H_i \cdot x_i^* + H_\delta \cdot \delta_j < 0)\ , \\ 0 & \text{if } (\sum_{i=1}^{m} H_i \cdot x_i^* + H_\delta \cdot \delta_j \geq 0)\ . \end{cases}$$

Figure 2 shows the results for solutions using $N = 362$ (red boxes) and $N = 100$ (blue boxes) scenarios. Each box represents the sampled probability distribution of the violation $V_N(x^*)$, eq. (3), displaying the maximum and minimum value, the mean, median and the σ confidence level. Note that, when using $N = 362$ scenarios, computed from eq. (9), we always remain below the prescribed violation level of $\epsilon = 10\%$ (red dashed line). Our analysis demonstrates that the violation level remains constant when the network size increases, even if the number of scenarios is fixed, as predicted in Del Duca et al. (2025). This observation suggests that our framework effectively

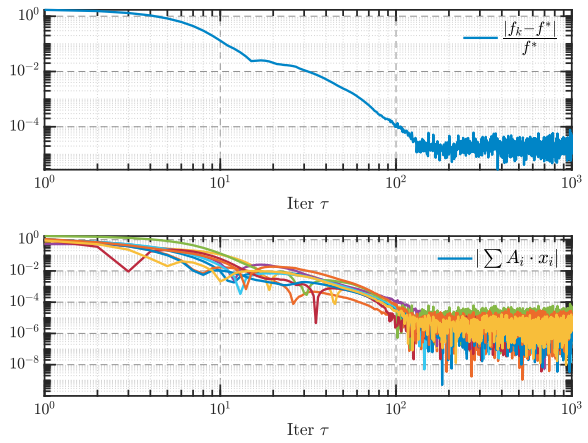

Fig. 3. Convergence of the normalised absolute error

separates the complexity of managing uncertainty from the overall complexity and structure of the network. As a result, the VPP can maintain operational efficiency without compromising the network's scalability.

Moreover, Fig. 2 shows that the violation levels for $N = 362$ are well below the prescribed violation level; we can make the solution less conservative by choosing a lower scenario number, for example, $N = 100$ (blue box); the resulting violation probability shifts toward the required risk level, while its variance increases, in line with the scenario theory from Campi and Garatti (2011).

5.2 Distributed algorithm convergence

Thanks to assumptions 2 and 3, the distributed solution converges to the solution of the centralised problem as the iterations pass. The upper plot of Fig. 3, showing the evolution of the normalised absolute error between the distributed solution $f_\tau(x_\tau)$ and the optimal centralised solution f^* and the lower plot, showing the residuals of the coupling equality constraints $e = P_i^j - P_j^i \ \forall (i,j) \in \mathcal{E}$, both exhibit that all the residuals converge to a reasonably small value ($\leq 10^{-4}$) in at most $\tau = 100$ iterations. As an additional note, the convergence of the errors depends mainly on the step size c in Alg. 4 and the connectivity of graph $\mathcal{G}$; in fact, the algorithm has a consensus dynamics on the dual variables, which depends on the graph connectivity and topology.

6. CONCLUSIONS

In this work, we presented a distributed architecture for managing a Virtual Power Plant that can simultaneously optimize a network of microgrids and provide the primary grid with a minimum level of aggregated power production while guaranteeing a-priori probabilistic constraint satisfaction levels on unseen net-demand realizations. Our method has several advantages, such as privacy, robustness to communication issues and disconnections, and scalability with respect to the number of agents. Due to its simplicity, efficiency, and scalability, the presented architecture is a step towards real-life implementations of distributed microgrid aggregating in a VPP. As forthcoming research, we are developing a VPP architecture tailored to solve MILP problems, increasing the fidelity of costs and dynamical models (e.g., batteries) while considering flexible loads and demand response.

REFERENCES

Baringo, A., Baringo, L., and Arroyo, J.M. (2019). Day-ahead self-scheduling of a virtual power plant in energy and reserve electricity markets under uncertainty. *IEEE Transactions on Power Systems*, 34(3), 1881–1894. doi: 10.1109/TPWRS.2018.2883753.

Campi, M.C. and Garatti, S. (2011). A Sampling-and-Discarding Approach to Chance-Constrained Optimization: Feasibility and Optimality. *Journal of Optimization Theory and Applications*, 148(2), 257–280. doi: 10.1007/s10957-010-9754-6.

Del Duca, A., Ruiz, F., and Scattolini, R. (2025). Distributed stochastic optimisation with uncertain coupling constraints. *IEEE 63rd Conference on Decision and Control (CDC)*.

Del Duca, A., Ruiz, F., and Scattolini, R. (2024). Robust micro-grid energy management system through a scenario approach. In *2024 American Control Conference (ACC)*, 1801–1806.

Diaz, C., Ruiz, F., and Patino, D. (2018). Smart charge of an electric vehicles station: A model predictive control approach. In *2018 IEEE Conference on Control Technology and Applications (CCTA)*.

Falsone, A. and Prandini, M. (2023). Augmented lagrangian tracking for distributed optimization with equality and inequality coupling constraints. *Automatica*, 157, 111269. doi: https://doi.org/10.1016/j.automatica.2023.111269.

Hannan, M.A., Abdolrasol, M.G.M., Faisal, M., Ker, P.J., Begum, R.A., and Hussain, A. (2019). Binary particle swarm optimization for scheduling mg integrated virtual power plant toward energy saving. *IEEE Access*, 7, 107937–107951. doi:10.1109/ACCESS.2019.2933010.

Liang, Z., Alsafasfeh, Q., Jin, T., Pourbabak, H., and Su, W. (2019). Risk-constrained optimal energy management for virtual power plants considering correlated demand response. *IEEE Transactions on Smart Grid*, 10(2), 1577–1587. doi:10.1109/TSG.2017.2773039.

Naughton, J., Wang, H., Cantoni, M., and Mancarella, P. (2021). Co-optimizing virtual power plant services under uncertainty: A robust scheduling and receding horizon dispatch approach. *IEEE Transactions on Power Systems*, 36(5), 3960–3972. doi: 10.1109/TPWRS.2021.3062582.

Vahedipour-Dahraie, M., Rashidizadeh-Kermani, H., Shafie-Khah, M., and Catalão, J.P.S. (2021). Risk-averse optimal energy and reserve scheduling for virtual power plants incorporating demand response programs. *IEEE Transactions on Smart Grid*, 12(2), 1405–1415. doi:10.1109/TSG.2020.3026971.

Vink, K., Ankyu, E., and Koyama, M. (2019). Multiyear microgrid data from a research building in Tsukuba, Japan. *Scientific Data*, 6(1). doi:10.1038/sdata.2019.20.

Yang, D., He, S., Wang, M., and Pandžić, H. (2020). Bidding strategy for virtual power plant considering the large-scale integrations of electric vehicles. *IEEE Transactions on Industry Applications*, 56(5), 5890–5900. doi:10.1109/TIA.2020.2993532.