Network-Secure Envelopes Enabling Reliable DER Bidding in Energy and Reserve Markets

Ahmad Attarha *Student Member, IEEE*, S. Mahdi Noori R.A., Paul Scott and Sylvie Thiébaux *Member, IEEE*

***Abstract*—As our power system modernises to embrace consumer-owned distributed energy resources (DER), network operators must ensure a safe and reliable operation while enabling consumers to trade their flexibility in the wholesale market. To enable this, we obtain network-secure operating envelopes that facilitate participation of consumers in energy and reserve markets. Given the distributed nature of the real-world problem, we use the alternating direction method of multipliers (ADMM) to optimise for operating envelopes such that any market action of consumers within these envelopes satisfies the distribution network constraints. To guarantee that the uncertainty realisation in live operation neither leads to network infeasibilities (due to exceeding the operating envelope) nor penalises DER-owners (due to market bid violation), we introduce a piecewise affinely adjustable robust bidding approach that can compensate for uncertainty variations in real-time. We also open up network capacity and minimise its losses, by proposing an additional piecewise affine Q-P controller that exchanges inverter reactive power with the grid. Our results on a 69-bus distribution network highlight the effectiveness of our proposal compared to alternative approaches.**

***Index Terms*—Distribution Network, Electricity Market, Operating Envelope, OPF, Reserve.**

I. Introduction

A. Research Motivation

The ongoing integration of distributed energy resources (DER), such as rooftop PV and batteries, has brought economic and environmental benefits to our power system. As a side effect however, the large generating units that traditionally provide frequency response are increasingly being pushed out of electricity markets. Aggregated DER has the potential to trade reserve services to compensate for this withdrawal. But this raises new challenges for the optimisation and control of our electricity system, which neither markets, networks, nor the demand side are currently prepared for.

One of these challenges is that the synchronisation of many DER, responding to price spikes, can exceed the technical limits of the distribution network sitting between the DER and the market. Another important issue is that the uncertainty about demand and PV generation can compromise the reliability of DER bids, and put our power system at a notable risk.

Our approach to these issues builds on the notion of *dynamic operating envelopes* [1]–[3], which has recently attracted the attention of system operators as a solution to ensure the security of the distribution network. An operating envelope is a convex set that defines the real and reactive power allowed to be transferred to / from the network at a given customer connection point, or for an aggregate of customers in a region. Envelopes are calculated so that any joint combination of consumption or generation within the envelope will not violate any network constraints.

Current proposals recommend that distribution system operators (DSOs) should repeatedly calculate operating envelopes and allocate them to each DER, without considering consumer preferences, uncertainty or market participation (e.g., [1], [2]). Since some consumers might need less flexibility (e.g., due to self consumption), such proposals can result in envelopes conservatively limiting market-participating consumers. More importantly, despite the initial motivation of guaranteeing network security, current proposals ignore the consumer-side of the envelopes, and do not check whether it is feasible for consumers to remain within their given envelopes; yet network constraints might be violated if they fail to do so, for instance due to uncertainty around solar PV generation. In addition, operating envelopes are not designed to guarantee reliable bidding in the electricity market. Therefore a new and comprehensive approach, closing these gaps, is essential.

In this paper, we enable consumers to participate in the wholesale market by obtaining operating envelopes that account for consumer preferences and uncertainty in a rigorous manner. To reflect the distributed nature of the problem and the possible privacy concerns of stakeholders, we decompose the problem into consumer and DSO subproblems and use the alternating direction method of multipliers (ADMM) to obtain the operating envelopes. More specifically, consumers iteratively negotiate with the DSO their preferred envelopes through ADMM, until convergence to consensus operating envelopes that satisfy consumer preferences and network limits. To ensure that consumers can commit to their envelopes and honour their bids to the market in live operation, we build a piecewise affinely adjustable robust constraint optimisation (PWA-ARCO) into our consumer subproblem. PWA-ARCO is an extension to the conventional affinely ARCO (AARCO) [4] which increases the flexibility by breaking the uncertainty set into more pieces, enabling the response to be better optimised when operating away from the worst case conditions.

B. Related Work

With the rising adoption of DER into the demand side, new regulations / solutions have been proposed to avoid overloading distribution networks. These solutions create an spectrum with central and local approaches at its two extremes. Central approaches, e.g., [5], optimise all consumers and the distribution network on a single platform to obtain bids

This work was done in the framework of the project "Optimal DER Scheduling for Frequency Stability". This project received funding from the Australian Renewable Energy Agency (ARENA), as part of ARENA's Advancing Renewables Program.

Authors are with the College of Engineering and Computer Science, The Australian National University (ANU), Canberra, Australia. Email: ahmad.attarha@anu.edu.au; mahdi.noori@anu.edu.au; paul.scott@anu.edu.au; sylvie.thiebaux@anu.edu.au

compliant with the grid technical limits. These approaches often lead to a large-scale optimisation problem which not only cannot scale to realistically sized networks, but also compromise consumer and network privacy. To overcome the challenges with a central tool, distributed optimisation ADMM is used in [6]–[9] to break the large-scale central problem into several smaller-scale subproblems which are coordinated using dynamic locational marginal prices (DLMPs). However, unlike [6]–[9], when bidding into the energy and reserve markets, there is no well-defined operating point, as the operating point depends on consumers' dispatch in the energy market and whether a contingency requiring a response occurs.

Contrary to the approaches based on either central [5] or distributed optimisation [6]–[9], local techniques limit consumer network access to a predefined value. Placing limits on power injected to the grid (e.g., on excess solar PV) is a common practice worldwide[1]. For example, in Germany, small-scale PV systems are not allowed to export more than 70% of their installed capacity [10]. Although simple, such fixed limits are overly conservative, as they are obtained for a scenario with maximum generation and minimum demand throughout a year.

To alleviate the overconservativeness of fixed limits, dynamic operating envelopes are proposed in [1]–[3]. Fig. 1 shows an example of an operating envelope – the green area represents the network secure operating points for the consumer. In [1] and [2] the DSO obtains operating envelopes while neglecting consumer preferences, uncertainty or market participation. Since some consumers might need less network access (e.g., due to self consumption), [1], [2] can result in envelopes conservatively limiting market-participating consumers. To obtain more precise envelopes, the DSO needs to have full observability over consumers [3] – i.e., detailed information around thousands of consumers that is not available to the DSOs. To conquer these challenges, inspired by [6]–[9], we use DLMPs within the distributed approach ADMM to enable consumers to negotiate their preferred operating envelopes with the DSO, rather than a single operating point as in [6]–[9]. This allows consumers to reflect their preferences into the optimisation, while providing some levels of privacy for consumers and the DSO. Notice that DLMPs can be used to create a new market within the distribution network. While this is a possible direction, in this paper we opt to minimise disruption to existing market structures, and so, once the network constraints are ensured, consumers participate in the main energy and reserve markets.

The network guarantees provided by envelopes hold as long as consumers can commit to their envelopes. However, envelopes (or any other deterministic approach such as central [5] / distributed [6]–[9], [11], [12] tools) are calculated prior to real time using forecast information. In live operation, the discrepancy between the forecast and reality might push consumers out of their envelopes, leading to a network violation. In addition, [2], [5]–[9], [11], [12] will pass on consumers' uncertainty to the electricity market, leading to bid deliveries different from the accepted offers. In the energy market, bid violations can be compensated for by regulation services which, in the Australian national electricity market (NEM), are paid for by the causers – a penalty not accounted for by the deterministic approaches. In reserve markets, not honouring bids can endanger power system's security. Thus, the NEM expels such participants from bidding into reserve markets, leading to a loss of a highly profitable revenue stream, again not accounted for by the deterministic approaches.

[1]When storage comes into the picture these limits need to be two-sided, i.e., both for power injection and power withdrawal scenarios.

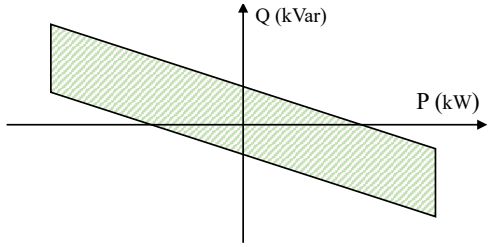

Fig. 1: An Example of an Operating Envelope [1]

To account for uncertainties, stochastic programming (SP) has been used, e.g., in demand response [13], network management [14], and unit commitment problem [15]. The scenarios used in SP can bias the solution away from the "true" PDF (probability distribution function). Thus, if the distribution of scenarios differs from the actual realisation, SP can result in poor out-of-sample performance. Moreover, considering a large set of scenarios to enhance the prediction of uncertainty realisation renders the problem computationally intractable. To overcome the challenges with SP, robust optimisation (RO) is used, e.g., in [16]–[21]. These works either neglect recourse capability, as in [16]–[19], and make decisions based on the worst case; or assume a linear recourse as in [20] and [21] that as we show in Section VI leads to over-conservative solutions.

Distributionally robust optimisation (DRO) has been used, e.g., in [22] and [23], to improve the performance of RO, based on the distribution of the uncertain parameters. However, the number of constraints of the DRO problem increases with the number of samples, leading to high dimensionality [24]. Hence, given the real-time market time frame (5 minutes in the NEM), these approaches are unlikely to scale, especially within the ADMM context where every subproblem needs to be solved several times for the algorithm to converge.

To meet the above challenges, we develop a piecewise affine framework that enables live correction in response to realisations that could belong to any PDF. Unlike [16]–[19], [22], [23] our approach can continuously tune the output based on what actually occurs in reality. In contrast to linear recourse [20], [21], we avoid over-conservative results by chunking the uncertainty set into pre-defined pieces and optimising for a piecewise linear function, each piece associated with a segment in the uncertainty set.

To ensure that consumers can commit to their envelopes, we integrate our piecewise affine functions into consumer side of operating envelopes. During the optimisation, the parameters of the piecewise functions are obtained such that consumers can stick to the envelope for any realisation within a polyhedral uncertainty set. To open up network capacity (i.e., wider operating envelopes), we enable consumers to additionally negotiate the reactive power support of their inverters with

the grid. This is in contrast to [11], [18], [25]–[27] that only rely on real power management. Moreover, we implement our approach within a model predictive control (MPC) framework which moves forward every 5 minutes in lock with the NEM real-time market. This allows us to feed the optimisation with the latest (most accurate) uncertainty information enabling the use of less conservative uncertainty sets. More information about the MPC framework is provided in Section V-C.

Compared to the DER coordination approach in [25], where consumers have to keep their CPP constant, here consumers CPPs can vary within an operating envelope. Moreover, consumers can here participate in both energy and reserve markets, leading to a higher value use of DER. Unlike the standard ARCO in [25], our PWA-ARCO produces a less conservative solution as it does not have to optimise only according to the worst-case scenario. To the best of our knowledge, we are the first to propose a PWA-ARCO solution approach for a consumer bidding problem.

Compared to [28] and [29] that obtain a flexibility region at the intersection of the transmission and distribution networks, we obtain envelopes down at consumer connection points while accounting for uncertainty. This enables consumers to provide their network-secure flexibility to the transmission network through the electricity market.

C. Contribution

In summary, we contribute to the state of the art by:

1) Proposing a distributed approach to obtain network-secure envelopes that account for consumer preferences and enable reliable market bidding. Unlike the operating envelopes presented in [1]–[3] that neglect consumer side, our approach uses the distributed optimisation ADMM to obtain envelopes that account for consumer preferences.

2) Proposing a piecewise affinely ARCO consumer bidding approach which is less conservative than conventional AARCO, leads to a higher value use of DER, and reduces the need for frequent ADMM negotiation. Contrary to the deterministic approaches [1]–[3], [5]–[9], [11], [12], our approach accounts for PV power and demand uncertainty. Unlike [16]–[19], [22], [23], our approach avoids over conservative results by enabling the response to adjust to realisations in live operation. In contrast to [20] and [21] that limit the recourse to a linear function, we optimise for piecewise affine functions that enable better response when operating away from the worst cases. Plus, different from [16]–[23], we characterise uncertainties within a distributed framework.

3) Proposing a Q-P controller that enables consumers to negotiate their reactive power support with the grid to increases the network throughput. Unlike [9], [11], [18], [25]–[27] that only count on real-power curtailment, our approach takes advantage of both real and reactive power of consumers. Unlike [21], our Q-P controller is piecewise, distributed, and works alongside our real-power bidding approach.

D. Assumptions

In this paper, we assume consumers participate in the energy and reserve markets within the NEM. Thus, compatible with the NEM, we co-participate in 7 real-time markets: 1 energy market, 3 contingency raise and 3 contingency lower reserve markets. These reserve markets (known as contingency FCAS markets) are differentiated based on their response time to 6-sec., 60-sec., and 5-min. raise / lower reserve markets. Nevertheless, both our network-secure framework and uncertainty characterisation are general and can be implemented for any electricity systems with high DER penetration.

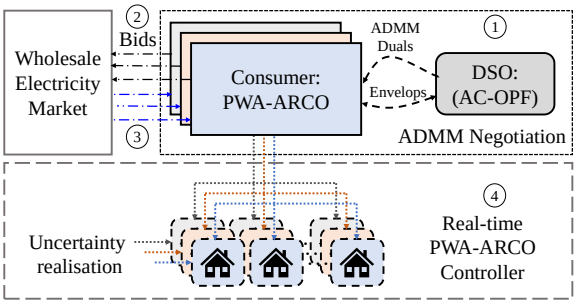

Fig. 2: High-level structure of PWA-ARCO.

Currently, small-scale participants are mainly able to participate in the wholesale market through a third party (e.g., an aggregator or a retailer). To avoid working with any specific number of aggregators / retailers, without loss of generality, we assume that each consumer itself can participate in the electricity market (i.e., each consumer is an aggregator / a retailer). However, our approach allows any number of aggregator / retailer.

E. Paper Organisation

The rest of this paper is organised as follows: Section II presents an upfront overall explanation on the proposed approach. Section III derives our PWA-ARCO methodology which is used in sections IV and V to develop our PWA-ARCO bidding approach. We provide the results and benchmarks against the state of the art in Section VI; finally, Section VII concludes this paper.

II. Overall Approach

Our main goal for every 5-minute interval is to obtain bids for the energy, 3 raise and 3 lower contingency reserve markets (7 markets in total) that are both robust against uncertainty and respect the grid operating limits. Fig. 2 shows a high-level scheme of our proposed approach which is repeated continuously using an MPC framework to generate and submit consumer bids to the wholesale market every 5 minutes. At the heart of our MPC framework lies an ADMM technique to enable every consumer and the DSO to negotiate every five minutes for an operating envelope that covers all market actions of the consumer. In order to robustify our bids against uncertainties and enable recourse actions, at every MPC iteration, consumers solve a PWA-ARCO problem to obtain their preferred market participation (bids) as well as the parameters of their piecewise functions. These piecewise functions will be enacted in real-time to deliver the bids according to the uncertainty realisation. We also include market bid deviation penalties in consumer objective function to help consumers make informed decisions about the parameters of

their piecewise functions. We also equip consumers with a Q-P controller which enables them to offer their reactive power support (generated / consumed by their inverters) to the DSO during the ADMM negotiations ①. At the convergence of our ADMM approach, the robust bids for each market as well as the parameters of our PWA-ARCO controllers are generated prior to realisation of uncertainty (the top right box in Fig. 1).

The network-secure bids are then submitted to the electricity markets ② which then get dispatched according to market clearing price ③. In live operation, the controllers take local recourse actions to keep consumers CPPs within the operating envelope upon which the network and consumers agreed at the convergence of ADMM, and avoid market penalties ④.

III. Building General Piecewise Affinely ARCO

We start with the conventional AARCO, upon which we build our proposed PWA-ARCO approach in Section III-B.

A. Affinely Adjustable Robust Constraint Optimisation

Let $\epsilon \in \mathbb{R}^l$ be the uncertain parameters, e.g., raise and lower reserve activation, solar power and demand uncertainty. We model ϵ in the following polyhedral uncertainty set:

$$E \triangleq \left\{ \epsilon \in \mathbb{R}^l | W\epsilon \leq v \quad : \mu \right\} \tag{1}$$

where $W \in \mathbb{R}^{k \times l}$ and $v \in \mathbb{R}^k$ are parameters of the polyhedral uncertainty set, and $\mu \in \mathbb{R}^k_{\geq 0}$ are dual variables. Let us begin with constraints of a robust optimisation problem in their most general form as:

$$\mathcal{B}x + \mathcal{C}\epsilon \leq d \qquad \forall \epsilon \in E \tag{2}$$

where $x \in \mathbb{R}^n$ is the vector of decision variables, $\mathcal{B} \in \mathbb{R}^{m \times n}$, $\mathcal{C} \in \mathbb{R}^{m \times l}$ and $d \in \mathbb{R}^m$. The for-all quantifier implies that (2) needs to be satisfied for any uncertainty realisation[2] within the uncertainty set E. To allow real-time recourse, an AARCO approach allows the decision variables to be an affine function of uncertainty as:

$$x \rightarrow x(\epsilon) \triangleq A\epsilon + b \tag{3}$$

where $A \in \mathbb{R}^{n \times l}$ and $b \in \mathbb{R}^n$. Prior to real-time we optimise to obtain A and b. In live operation, when the true value of ϵ is revealed, $x(\epsilon)$ will become fully known. This is in the spirit of a linear feedback controller, where the value of x can be constantly updated in response to the realisation of uncertainty.

We next substitute (3) into (2) and rewrite (2) in its equivalent robust form using a max protection function on a per-constraint basis [4]. Using $\mathcal{B}_i$, $\mathcal{C}_i$ and d_i to represent the i-th row of their respective matrices / vectors, we get:

$$\mathcal{B}_i b + \max_{\epsilon \in E} \left(\mathcal{B}_i A\epsilon + \mathcal{C}_i \epsilon \right) \leq d_i \tag{4}$$

This maximisation problem is an LP for any given $\mathcal{B}$, A and $\mathcal{C}$; E is a polyhedron and the objective is linear in ϵ, so, we can replace the max inside (4) with its dual form [4] as:

$$\mathcal{B}_i b + v^\intercal \mu_i \leq d_i \tag{5a}$$

$$W^\intercal \mu_i \geq \mathcal{B}_i A + \mathcal{C}_i \tag{5b}$$

[2]Notice that in the deterministic case (2) needs to be satisfied only for the forecast scenario.

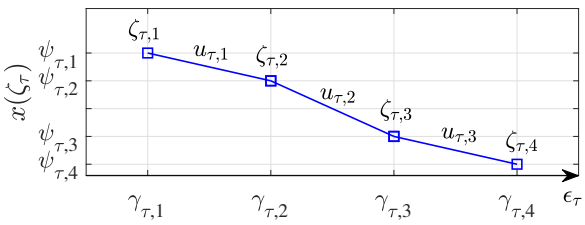

Fig. 3: Visualisation of a piecewise affine function

where $\mu_i \in \mathbb{R}^k_{\geq 0}$ are the dual variables associated with (1) and (5a)–(5b) represent the constraint-wise robust counterpart of problem (4). Constraints (5a)–(5b) ensure that constraint (4) can be satisfied for any realisation of uncertainty within the uncertainty set E (i.e., the affine function $x(\epsilon)$ can be adjusted to compensate for all uncertainty realisations within E). Notice that here we only focus on the constraints. The epigraph form of the objective function can be written and thus the robust objective can also be modelled as in (2), alternatively, the objective may weigh up the outcome under a selection of scenarios.

Immunising constraint (4) via affine functions can lead to over conservative solutions, as it forces to respond with the same slope (A in $A\epsilon + b$) for any realisation within the uncertainty set. To enable consumers to tune their behaviour for different uncertainty realisation, we introduce the PWA-ARCO in the following.

B. Piecewise Affinely ARCO

The idea here is to partition the uncertainty set into contiguous subsets, and then optimise x(·) as a piecewise affine function. We begin with encoding our piecewise affine functions (see Fig. 6 for a visualisation of a piecewise affine function with 4 breakpoints):

$$x(\zeta) \triangleq \sum_{\tau=1}^{l} \psi_\tau \zeta_\tau \tag{6}$$

Where $\psi_\tau \in \mathbb{R}^{n \times \alpha_\tau}$ are the "y-axis" breakpoints for each variable, for the τ-th random parameter (what we optimise), α_τ is the number of breakpoints and $\zeta_\tau \in [0,1]^{\alpha_\tau}$ are weights for the τ-th random parameter. These weights encode the position in the uncertainty set through:

$$\epsilon_\tau = \gamma_\tau^\intercal \zeta_\tau \qquad \sum_{a=1}^{\alpha_\tau} \zeta_{\tau,a} = 1 \tag{7}$$

Where $\gamma_\tau \in \mathbb{R}^{\alpha_\tau}$ are the "x-axis" breakpoints (each entry is unique) for the τ-th random parameter. Notice that the weights for each random parameter must sum to 1.

To further restrict the weights, so they encode a piecewise linear function, we introduce binary variables $u_\tau \in \{0,1\}^{\alpha_\tau - 1}$ and the following set of constraints:

$$\sum_{a=1}^{\alpha_{\tau-1}} u_{\tau,a} = 1 \tag{8a}$$

$$\zeta_{\tau,a} \leq u_{\tau,a-1} + u_{\tau,a} \qquad \forall a \in \{2, ..., \alpha_\tau - 1\} \tag{8b}$$

$$\zeta_{\tau,1} \leq u_{\tau,1} \tag{8c}$$

$$\zeta_{\tau,\alpha_\tau} \leq u_{\tau,\alpha_\tau - 1} \tag{8d}$$

When ensuring (2) holds for all values in the uncertainty set, using a max protection function per constraint i, we get:

$$\max_{\epsilon \in E} \left(\mathcal{B}_i x(\epsilon) + \mathcal{C}_i \epsilon\right) \leq d_i \tag{9}$$

When constraints (7) and (8a)–(8d) hold, there is a unique mapping from ϵ to ζ, allowing us to replace $x(\epsilon)$ with $x(\zeta)$. Expanding out the maximisation we have:

$$\max_{\zeta,\epsilon} \left(\mathcal{B}_i \sum_{\tau=1}^{l} \psi_\tau \zeta_\tau + \mathcal{C}_i \epsilon\right) \tag{10a}$$

$$W\epsilon \leq v \tag{10b}$$

$$(7),\ (8a)\text{–}(8d) \tag{10c}$$

Equations (10a)–(10c) represent an MILP problem and thus it is not possible to directly use duality theory and reformulate it, so, a similar approach as in (5a)–(5b) will not work here. Notice that if we relax the binaries, the solution will be greater or equal to the LHS of (9). This means that if we utilise the binary relaxed formulation in order to obtain a dual, our solution will still be robust, but potentially more than it needs to be. Therefore, such a relaxation does not compromise feasibility as all of the constraints are still guaranteed. In the result section, we compare our results with a perfect case and discuss the conservativeness of different approaches. Relaxing the binary variables u_τ allows $(\epsilon, \psi_\tau \zeta_\tau)$ to take on a position anywhere in the convex hull of the breakpoints of the relaxed piecewise linear function (recall $\zeta_{\tau,a} \in [0,1]$). Writing the relaxation out explicitly, and indicating how the duals associate with the constraints, we have:

$$\max_{\epsilon,\zeta} \left(\mathcal{B}_i \sum_{\tau=1}^{l} \psi_\tau \zeta_\tau + \mathcal{C}_i \epsilon\right) \leq d_i \tag{11a}$$

$$W\epsilon \leq v \qquad (\mu \geq 0) \tag{11b}$$

$$\epsilon_\tau - \gamma_\tau^{\intercal} \zeta_\tau = 0 \qquad (\mu'_\tau) \tag{11c}$$

$$\sum_{a=1}^{\alpha_\tau} \zeta_{\tau,a} = 1 \qquad (\mu''_\tau) \tag{11d}$$

$$\zeta_{\tau,a} \leq 1 \qquad (\mu'''_{\tau,a} \geq 0) \tag{11e}$$

where μ, μ', μ'', and μ''' are dual variables. Taking the dual leads to:

$$\min_{\mu - \mu'''} \left(v^{\intercal}\mu + \sum_{\tau=1}^{l} \mu''_\tau + \sum_{\tau=1}^{l}\sum_{a=1}^{\alpha_\tau} \mu'''_{\tau,a}\right) \leq d_i \tag{12a}$$

$$W^{\intercal}\mu + \mu' \geq \mathcal{C}_i \tag{12b}$$

$$\mu''_\tau - \gamma_{\tau,a}\mu'_{\tau,a} + \mu'''_{\tau,a} \geq \mathcal{B}_i \psi_{\tau,a} \tag{12c}$$

Recall that these dual variables and constraints will have a copy for each constraint in the original problem (i.e. a further index of i that has been treated implicitly in the above). Notice that if (12a) is satisfied for a value of $\left(v^{\intercal}\mu + \sum_{\tau=1}^{l} \mu''_\tau + \sum_{\tau=1}^{l}\sum_{a=1}^{\alpha_\tau} \mu'''_{\tau,a}\right)$, then it will be satisfied for its minimum as well. Therefore we can drop the $\min$ operator. The resulting problem (12a)–(12c) represents our PWA-ARCO solution approach.

IV. High-Level Consumer and Network Negotiation

In this section, we aim to provide a high-level presentation of the ADMM approach and what is being negotiated between consumer and the DSO. In Section IV-A, we only explain the uncertainty sets and our piecewise affine operating envelopes based on the theory we just covered. We then present our high-level network subproblem and the ADMM approach in Section IV-B. The details of the consumer objective and constraints that interact with the envelope as well as DSO constraints are left until Section V.

A. Consumer Subproblem

In our bidding problem, the uncertain parameter ϵ includes FCAS activation $\epsilon^F \in [-1, 1]$, PV power $\epsilon^{PV} \in [\overline{P}^{PV}, \underline{P}^{PV}]$ and residential demand $\epsilon^D \in [\overline{P}^D, \underline{P}^D]$. Regarding the FCAS activation, notice that raise and lower FCAS services will not be activated at the same time, plus, the activation of 6-sec, 60-sec and 5-min reserves (raise or lower) are sequential, e.g., if 5MW is bid into 6-sec, 60-sec and 5-min, then at most 5MW response is required at any one moment (rather than 15MW). Notice that network constraints need to be satisfied independently of the type of raise or lower reserve response. Therefore, we use a single uncertain parameter ϵ^F to account for the most extreme reserve activation across all markets. Similarly to $x(\zeta)$, we write the real power generated by a consumer $p(\zeta) \in \mathbb{R}^{|T|}$ as:

$$p(\zeta) \triangleq \sum_{a=1}^{\alpha_F} \psi_a^F \zeta_a^F + \sum_{a=1}^{\alpha_{PV}} \psi_a^{PV} \zeta_a^{PV} + \sum_{a=1}^{\alpha_D} \psi_a^D \zeta_a^D \tag{13a}$$

$$\epsilon^F = \gamma^{F\intercal} \zeta^F \qquad \sum_{a=1}^{\alpha_F} \zeta_a^F = 1 \tag{13b}$$

$$\epsilon^{PV} = \gamma^{PV\intercal} \zeta^{PV} \qquad \sum_{a=1}^{\alpha_{PV}} \zeta_a^{PV} = 1 \tag{13c}$$

$$\epsilon^D = \gamma^{D\intercal} \zeta^D \qquad \sum_{a=1}^{\alpha_D} \zeta_a^D = 1 \tag{13d}$$

(13a)–(13d) correspond to (6)–(7) yet with the uncertain parameters broken up into independent subsets. Notice that $p(\zeta)$ represents all possible CPPs of the consumer for any realisation of uncertain parameters ϵ. To increase readability, we interchangeably use $p(\zeta)$ and $p(\epsilon)$ throughout this paper (both meaning that p is a function of uncertain parameters). To have an upper and lower bound on these CPPs, we define the vector variables $\underline{p}, \overline{p} \in \mathbb{R}^t$ indicating the operating envelope where:

$$p(\epsilon) \in [\underline{p},\ \overline{p}] \qquad \forall \epsilon \in E \tag{13e}$$

The above equation indicates that the CPP resulting from participating in any or all 7 energy and reserve markets for any realisation of uncertainty within the polyhedron E will remain within the operating envelope $[\underline{p},\ \overline{p}]$. When network constraints are neglected, the operating envelope will not limit the market actions, i.e., $[-\infty,\ \infty]$. However, this envelope might significantly limit consumers when the network constraints are taken into account. To increase network throughput and to enable

consumers to bid with less network restriction, we obtain the reactive power support associated with the operating envelope with the following Q-P controller:

$$q(\zeta) \triangleq \sum_{a=1}^{\alpha_F} \psi_a^{'F} \zeta_a^F + \sum_{a=1}^{\alpha_{PV}} \psi_a^{'PV} \zeta_a^{PV} + \sum_{a=1}^{\alpha_D} \psi_a^{'D} \zeta_a^D \quad (13f)$$

We use the above decision rule to obtain the reactive power that can be exchanged with the grid at the uncertainty realisation ζ. Let $q_{\underline{p}}$ be the reactive power associated with $\underline{p}$ and $q_{\overline{p}}$ denote the reactive power associated with $\overline{p}$. When reactive power is neglected both $q_{\underline{p}}$ and $q_{\overline{p}}$ are zero. More details about this is provided in Section V. Next, we negotiate the pairs $\underline{s} \triangleq (\underline{p}, q_{\underline{p}})$ and $\overline{s} \triangleq (\overline{p}, q_{\overline{p}})$, output of consumer energy management system (EMS), with the grid. The DSO then solve two OPFs one for each pair to ensure network feasibility. To ease the presentation, we put $\underline{s}$ and $\overline{s}$ in a vector s, where $s \triangleq (\underline{s}, \overline{s})$.

B. Network Subproblem and ADMM Algorithm

So far, at a high level, we explained the operating envelope s obtained at the consumer subproblem. To ensure that consumer envelopes are network secure, here we negotiate them with the DSO using the ADMM approach. This completes (1) in Fig. 1. Using s' for the same variable (i.e., duplication of s) but from the network perspective, in the final solution we have:

$$s - s' = 0 \qquad (\lambda) \quad (14a)$$

We write the augmented Lagrangian associated with (14a) as:

$$\mathcal{L}^*(s, s', \lambda) = \lambda^\intercal (s - s') + \frac{\rho}{2} ||s - s'||_2^2 \quad (14b)$$

where λ is a vector of dual variables of (14a) and ρ is the penalty parameter of the augmented Lagrangian. We use the ADMM algorithm [30] as a negotiation tool between consumers and the network to obtain the network-secure results. Given the consumer objective function $f(.)$, the network internal variables y, and the objective and constraint functions $\mathcal{G}$ and $\mathcal{H}$, our ADMM approach solves the following per iteration k:

$$s^{(k)} := \min_{x(\epsilon)} \left[f(.) + \mathcal{L}^*(s, s'^{(k-1)}, \lambda^{(k-1)}) \right]$$
$$\text{s.t.} \qquad (12a)\text{–}(12c),\ (13a)\text{–}(13f) \quad (15a)$$
$$s'^{(k)} := \min_{z} \left[\mathcal{G}(s', y) + \mathcal{L}^*(s^{(k)}, s', \lambda^{(k-1)}) \right]$$
$$\mathcal{H}(s', y) \leq 0 \quad (15b)$$
$$\lambda^{(k)} := \lambda^{(k-1)} + \rho^{(k)} (s^{(k)} - s'^{(k)}) \quad (15c)$$

In the first phase (15a), consumers are optimised for s, while holding s' and λ constant at their $k-1$-th value. In the second phase (15b), the network is optimised for s', while holding s and λ constant at their k and $k-1$-th values respectively. Finally, the dual variables λ are updated in (15c), completing the k-th iteration. In line with [11], our approach converges if the infinity norms of the primal $||R_p^{(k)}||_\infty$ and dual residuals $||R_d^{(k)}||_\infty$ are both smaller than a threshold.

Remark 1: Here, we only solve two OPFs associated with the most extreme operating points, i.e., $\underline{s}$ and $\overline{s}$, in which consumers simultaneously inject / absorb power to / from the grid. Since no network constraint is violated at these worst cases, we argue that the network constraints are also satisfied for less extreme scenarios in between. It is worth acknowledging that this assumption holds for radial networks that operate within the voltage stable mode. We provide a proof for this in Appendix A. Since most real-world distribution networks are radial (or operated radially), we believe the radial-network assumption does not limit the application of our approach in a real-world setting. We leave the extension of operating envelopes for mesh networks to future work.

Remark 2: The dual values λ represent the DLMPs associated with their respective envelopes. In this paper, we use dual variables to signal network constraints and ensure reliable and robust consumer bidding in the wholesale market. We acknowledge that at the last step dual values can be used to charge customers for their network access. However, this opens up non-trivial questions such as: how these prices should interact with the current fixed network access fees, or whether additional DSO income (on top of fixed access fees) is justified. Plus, DSOs might need schemes to ensure that customers, located at the end of the feeders, are not overly-disadvantaged by DLMPs. Answering these questions are out of the scope of the current study. We leave further investigations into this aspect to future work.

V. Detailed Consumer and Network Subproblem

Here, we develop a PWA-ARCO energy management system (EMS) problem for a consumer participating in energy, and 6-sec, 60-sec, and 5-min raise and lower reserve markets which is then followed by a detailed DSO subproblem.

A. Detailed EMS Subproblem

1) DER Constraints: We need to ensure the DER constraints are satisfied for any CPP realisation, resulting from reserve market activation or PV and demand realisations. To do so, we represent DER variables to be piecewise affine functions of uncertainty and then use the piecewise affine function to write each constraint. The consumer subproblem also includes real-power (13a), the uncertainty set (13b)–(13d), the envelope (13e) and the reactive power (13f).

Solar PV: Solar PV forecast has been modelled in the uncertainty set. We also use a piecewise affine function $p_t^{cur}(\zeta_\tau) \in [0, \epsilon_t^{PV}]$ to model the curtailment as a function of uncertainty. Using the max protection function we have:

$$\max_{\zeta_\tau} \{ p_t^{cur}(\zeta_\tau) - \gamma_t^{PV\intercal} \zeta_t^{PV} \} \leq 0 \quad (16a)$$

Battery Storage: We define the piecewise affine functions for battery charge $p_t^C(\zeta_\tau) \in [0, \mathcal{R}]$ and discharge $p_t^D(\zeta_\tau) \in [0, \mathcal{R}]$ variables. Given battery efficiency η, the protection functions for battery's bounding constraints can be written as:

$$\min_{\zeta_\tau} \{ p_t^C(\zeta_\tau) \} \geq 0; \qquad \max_{\zeta_\tau} \{ p_t^C(\zeta_\tau) \} \leq \mathcal{R} \quad (16b)$$
$$\min_{\zeta_\tau} \{ p_t^D(\zeta_\tau) \} \geq 0; \qquad \max_{\zeta_\tau} \{ p_t^D(\zeta_\tau) \} \leq \mathcal{R} \quad (16c)$$
$$\min_{\zeta_\tau} \{ e_0 + \sum\nolimits_{\delta=1}^{t} (\eta p_\delta^C(\zeta_\tau) - p_\delta^D(\zeta_\tau)/\eta) \} \geq \underline{e} \quad (16d)$$
$$\max_{\zeta_\tau} \{ e_0 + \sum\nolimits_{\delta=1}^{t} (\eta p_\delta^C(\zeta_\tau) - p_\delta^D(\zeta_\tau)/\eta) \} \leq \overline{e} \quad (16e)$$

where $\mathcal{R}$ is the charging / discharging rate and $\underline{e}$ and $\overline{e}$ are the minimum and maximum values for the battery SoC. To ensure that simultaneous charge and discharge does not occur, we use a binary variable u_t and the following constraints:

$$0 \leq p_t^C(\overline{\zeta}) \leq \mathcal{R} \cdot u_t \tag{16f}$$
$$0 \leq p_t^C(\overline{\zeta}) \leq \mathcal{R} \cdot (1 - u_t) \tag{16g}$$

We only apply these binaries to the energy market, i.e., $\zeta = \overline{\zeta}$. Because in reserve markets, the battery can transit form charge to discharge (or visa versa) to provide a greater response.

Combined Power: The combined household power can be written as follows:

$$p(\zeta_\tau) = p_t^D(\zeta_\tau) - p_t^C(\zeta_\tau) - p_t^{cur}(\zeta_\tau) + \epsilon_t^{PV} + \epsilon_t^D \tag{16h}$$

The equality constraint (16h) gives the relation between the combination of the piecewise affine parameters on the LHS with the piecewise affine parameters of the CPP on the RHS of (16h). Having obtained all the protection functions, we use duality theory to convert each protection functions into some linear constraints as in (12a)–(12c).

2) Objective Function: The market payments to the customer consist of what energy they exchange with the network and the reserve market commitments (whether deployed or not). They are also penalised if the energy they exchange deviates from their energy market amount, unless this deviation is accounted for by reserve market activation. What we propose to do is evaluate the revenue from the energy market at the most likely realisation of uncertainty (i.e., forecast, shown by $\bar{\epsilon} = \{\epsilon^F = 0, \epsilon^{PV} = \bar{\epsilon}^{PV}, \epsilon^D = \bar{\epsilon}^D\}$, given by $p(\bar{\epsilon})$. This will also be the amount we bid into the energy market. Notice that this is not equivalent to making energy bids deterministically as the output of our functions also account for uncertainties at the forecast scenarios $\bar{\epsilon}$.

For the reserve markets, we assume that we have to meet the bid capacity under all circumstances. Also, the reserve activation ϵ^F can vary between -1 (max lower activation) to 1 (max raise activation) where $\epsilon^F = 0$ represents a case where no reserve is required (energy case). Using r' for raise and l' for lower, we can write:

$$r' = p(\epsilon^F = 1, \epsilon^{PV}, \epsilon^D) - p(\epsilon^F = 0, \epsilon^{PV}, \epsilon^D) \tag{17a}$$
$$l' = p(\epsilon^F = 0, \epsilon^{PV}, \epsilon^D) - p(\epsilon^F = -1, \epsilon^{PV}, \epsilon^D) \tag{17b}$$

Notice that the contribution of ϵ^F in $p(\epsilon)$ is separable from the other sources of uncertainty, because it is an independent piecewise affine function and ϵ^F can be considered not part of the E polyhedron. This means the rest (uncertain PV and demand) cancel out, and we do not have to resort to doing a forall ϵ. In other words, our reserve bids are robust to PV and demand uncertainty and are deliverable for any realisations within the PV and demand uncertainty sets. This is an important feature of our approach since it makes consumers reliable reserve providers. In the NEM, if a participant cannot honour its reserve bids (e.g., a likely scenario for consumers due to uncertainty), they will be excluded from future reserve market participation. However, by obtaining deliverable reserve bids, our approach secures a spot for consumers in these highly prices yet rarely activated markets.

Since contingency reserve markets in the NEM are activated sequentially, the NEM allows the same capacity to be submitted to any or all contingency reserve markets. Let r^6 / l^6, r^{60} / l^{60}, and r^5 / l^5 respectively denote raise / lower reserve offers to 6-sec, 60-sec, and 5-min reserve markets, using (17a) and (17b), we have:

$$r^6 \leq r'; \qquad r^{60} \leq r'; \qquad r^5 \leq r' \tag{17c}$$
$$l^6 \leq l'; \qquad l^{60} \leq l'; \qquad l^5 \leq l' \tag{17d}$$

Let π^e denote the energy market price; $\pi^r = \{\pi^{r6}, \pi^{r60}, \pi^{r5}\}$ and $\pi^l = \{\pi^{l6}, \pi^{l60}, \pi^{l5}\}$ be the raise and lower 6-sec, 60-sec and 5-min market prices, $r = \{r^6, r^{60}, r^5\}$ and $l = \{l^6, l^{60}, l^5\}$ be the offers to the raise and lower reserve market. The obtained benefit can be modelled as follows:

$$C^{Ben} = \pi^e p(\bar{\epsilon}) + \pi^r r + \pi^l l \tag{17e}$$

The above equation calculates the benefits obtained in the energy and contingency reserve market. However, as we mentioned earlier there is a penalty associated with bid violation in energy market. To model this penalty, we interpret the causer pays policy in the NEM, to be penalising energy-bid deviation at the regulation market price[3]. We add the penalty C^{Pen} to the objective function to provide consumers with information about the bid deviation penalty, helping them to make more informed decisions. Given the penalty price π^- this can be written as:

$$C^{Pen} = \pi^- |p(\bar{\epsilon}^F, \bar{\epsilon}^{PV}, \bar{\epsilon}^D) - p(\bar{\epsilon}^F, \epsilon^{PV}, \epsilon^D)| \tag{17f}$$

we use auxiliary variables to model the absolute value function $|\cdot|$ in a linear manner, which is then treated robustly.

3) Reactive Power Network Support: The reactive power generated / consumed by the inverter can be modelled as:

$$q^2 \leq s^2 - {p^{inv}}^2 \tag{18a}$$

The above equation is a circle in (q, p^{inv}) coordinates, which can be linearised using a set of linear constraints as follows:

$$q(cos(\phi) + sin(\phi)) \leq \sqrt{2}s - p^{inv}(cos(\phi) - sin(\phi)) \tag{18b}$$

where $\phi \in \{0, \pi/b, 2\pi/b, ..., (2b-1)\pi/b\}$, and b is an arbitrary integer number. Here, we use 24 lines [4] (i.e., an icositetragon) which overestimate the circle with at most 0.001% error. Note that p^{inv} in our case can be battery charge / discharge and PV power which are modelled via piecewise affine functions $x(\zeta)$. As explained in (13f) q is also a function of uncertain parameters. Using $x(\zeta)$ and $q(\zeta)$, (18b) can be written as:

$$q(\zeta)(cos(\phi) + sin(\phi)) \leq \sqrt{2}s - x(\zeta)(cos(\phi) - sin(\phi)) \tag{18c}$$

The goal to provide reactive power support is to improve voltages so that consumers can have a bigger operating envelope. However, grid and consumers only negotiate on the required reactive power at the extremes. Thus, this does not provide any information about the required reactive power

[3]The reason for this is that regulation market is activated to compensate for violations in the energy market.

[4]Due to the distributed nature of our approach, each consumer problem can be solved separately. Thus, adding even more lines does not significantly add to the computational complexity.

away from the extremes. To obtain the required reactive power within the envelope, we take advantage of the results presented in [21], where the authors provide a closed-form solution to the problem of how much reactive power injection / absorption is needed at a node to fully compensate voltage changes due to real power deviations at that node. As discussed in [21], a $q(\zeta)$ affine function with a slope equal to the closed-form solution R^j (where j is a network node) would completely compensate the voltage deviation due to the active power deviations. This, for τ-th uncertain parameter, can be written as:

$$\frac{\partial q(\zeta)}{\partial \zeta_\tau} \leq R^j \frac{\partial p(\zeta)}{\partial \zeta_\tau} \tag{18d}$$

Finally, given the piecewise function (5) the partial differential equations can be written as:

$$\psi'_{\tau,i+1} - \psi'_{\tau,i} \leq R^j \psi_{\tau,i+1} - R^j \psi_{\tau,i} \quad \forall \tau, i \in \{1, ..., \alpha_\tau\} \tag{18e}$$

where R^j is the closed-form solution from [21]:

$$R^j = -\frac{\sum_{j=2}^{n} K_{jk}^{|v|q} K_{jk}^{|v|p}}{\sum_{j=2}^{n} K_{jk}^{|v|q^2}} \tag{18f}$$

where, $K_{jk}^{|v|q} = \frac{\partial v_j}{\partial q_k}$ and $K_{jk}^{|v|p} = \frac{\partial v_j}{\partial p_k}$. In summary, the constraints of our piecewise affinely ARCO EMS subproblem for a single customer consist of the objective function (17a)–(17f), the affine functions (13a)–(13f), DER operating constraints (16a)–(16h), and the Q-P controller (18c)–(18e).

B. Network Subproblem

The network subproblem needs to solve two multi-period OPFs associated with the extremes of s' at every MPC iterations. Notice that our ADMM approach decomposes consumers and the network. Thus, there are no time coupling constraints in the network subproblem, because all such constraints are part of the consumer subproblems, e.g., the SoC coupling constraint (16d) and (16e). This means that the OPF for each time step can be solved separately and in parallel.

In the following, we use the Dist-flow equations to model our network subproblem for the upper bound of $\overline{s}'$, i.e., $\overline{p}'_n$ and $q'_{\overline{p}_n}$. Index t is dropped to increase readability. We duplicate similar variables and constraints for $\underline{s}'$. We use $i, j, k \in N$ for nodes in a tree network; F_i^{ac}, F_i^{re} and I_i are the active power, reactive power and the current flowing into node i from the parent node k, where the line has resistance r_i, reactance x_i and impedance z_i. D_i represents the child nodes of node i. The network subproblem for $\overline{s}'$ can be written as follows:

$$\{\overline{p}'_n \; q'_{\overline{p}_n}\} := \underset{\overline{p}' \; q'_{\overline{p}}}{\operatorname{argmin}} \sum_{i \in N} \left[\pi^e r_i I_i + \mathcal{L}^*(.)\right] \tag{19a}$$

$$F_i^{ac} - r_i I_i - \sum_{n \in C^i} \overline{p}'_n = \sum_{j \in D_i} F_j^{ac} \quad \forall i \in N, \tag{19b}$$

$$F_i^{re} - x_i I_i - \sum_{n \in C^i} q'_{\overline{q}_n} = \sum_{j \in D_i} F_j^{re} \quad \forall i \in N, \tag{19c}$$

$$V_i = V_k - 2\Big(r_i F_i^{ac} + x_i F_i^{re}\Big) + z_i^2 I_i \quad \forall i \in N \tag{19d}$$

$$v_{min}^2 \leq V_i \leq v_{max}^2 \quad \forall i \in N \tag{19e}$$

$$F_i^{ac2} + F_i^{re2} = V_i I_i \quad \forall i \in N \tag{19f}$$

$$0 \leq I_i \leq i_i^{max2} \quad \forall i \in N \tag{19g}$$

The objective function is to minimise the total loss as well as the augmented Lagrangian penalty function $\mathcal{L}^*(.)$. Active and reactive power flow equations are given through (19b)–(19c); The voltage of each node is calculated through (19d) and is enforced to be within its safe limits (v_{min}^2 and v_{max}^2) through (19e). The complex power, flowing in each line, is given in (19f) and finally, (19g) limits the current of each line to the maximum line capacity i_i^{max2}. Note that due to having generation (PV and battery), a convex relaxation of (19f) might not lie within the feasible region of the original problem. Thus, to avoid infeasible results, we opt to solve the exact non-convex OPF using the IPOPT solver. Similarly to [8], [9], [11], we have also observed that our original OPF model could be solved efficiently by the IPOPT solver within the ADMM context. Nevertheless, we acknowledge the need for further investigations around theoretical convergence proof.

C. Model Predictive Control Implementation

The length of our MPC window is 24 hours discretised every 5 minute, i.e., 288 time steps per day. Every 5 minutes, in lock with the NEM real-time market, consumers and the DSO optimise to obtain envelopes, energy-reserve bids, and the parameters of the piecewise affine controllers. If during a 5-minute operating interval a reserve response gets called, our MPC accounts for it in the next optimisation by using the latest battery SoC.

We presented our PWA-ARCO treatment for all t in Section IV. However, when integrated within our MPC framework, we apply the robust treatment for PV power and demand only to the first time step. The reason is that our MPC framework enables us to use the latest (most accurate) forecast information every 5 minute when a new optimisation problem is solved. Since the uncertainty variation within the next 5 minute is often insignificant, the uncertainty set can cover the deviations from the forecast. On the contrary, constructing the uncertainty set for future time steps might lead to over-conservative results, as the forecast values (around which we construct the uncertainty set) can change significantly. In our simulations, this could bring 3% higher benefits compared to the case where the entire horizon was treated robustly.

VI. Numerical Results

To illustrate the effectiveness of the proposed approach, we modify a 69-bus distribution network [11] with 3 consumers at each node. Consumers connecting to the grid are consuming 20 kWh per day on average and own a 5 kW rooftop PV and a 5kW / 10kWh battery with the round trip efficiency of 81%. We use anonymised solar and demand data of consumers in Tasmania, Australia [9], and randomly assign this data to 207 consumers in our 69-bus network. The 5-minute energy and reserve market prices for the day 24/01/2020 are taken from the NEM website. Every 5 minute, we measure PV power and demand and use this values as forecasts. We then construct our uncertainty set to cover uncertainty realisations within $\pm 20\%$ of the forecast values. Five breakpoints is used to chunk the uncertainty sets resulting in four equal segments in our piecewise functions. The ADMM penalty parameter ρ for active and reactive power is chosen to be 1 [11] and

0.01, respectively[5]. Finally, we use the Gurobi and IPOPT solvers in JuMP, Julia [31] to solve our consumer and network subproblems, respectively.

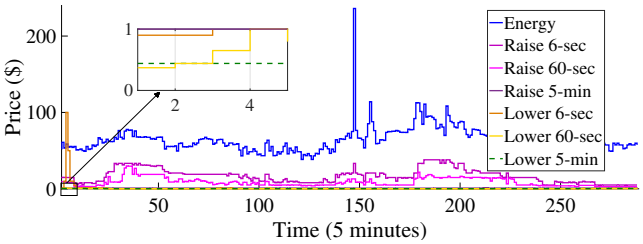

Fig. 4: Energy and FCAS prices

To assess the performance of our proposed PWA-ARCO approach, we study the following 4 approaches:

Deterministic: This is the work [11] in which consumers and the grid negotiate every five minutes through the standard multi-period ADMM algorithm (15a)−(15c). The consumer subproblem is only satisfied for the forecast scenario.

Perfect: This approach intends to provide a perfect but unachievable baseline in which (assuming that computation and communication time is not an issue) two multi-period OPFs (associated with the operating envelope) are solved every minute to ensure the grid feasibility for all scenarios. Thus, neither the piecewise controllers nor the binary relaxation does apply to this case. Since all the future information is assumed to be available, the FCAS bid offers of this case are all made to be deliverable.

AARCO: This is the conventional affinely ARCO where we use the affine decision rule $x(\epsilon) \triangleq A\epsilon + b$ and (5a)–(5b) to obtain our robust bids.

PWA-ARCO: This is the proposed approach in which the grid envelopes and the parameters of our piecewise affine controllers are obtained using a grid-wide ADMM coordination. In real-time our controllers are continually taking recourse actions to keep the CPP within the negotiated envelope.

We implement the above approaches within an MPC framework which moves forward every 5 minutes. During every 5-minute operating interval, we use the 1-minute data as live realisations, to evaluate the effectiveness and real cost of each approach. Having the whole horizon covered, we track the SoC of batteries and move to the next horizon in which the same process is repeated until the whole horizon is covered, i.e., 288 horizons with 1440 realisation scenarios. In the following, we first study the 1-day performance of our approach in details and then provide a 3-month experiment to evaluate the longer-term performance of our approach.

A. Total Benefit in Energy and Reserve Markets

Table I reports the total benefits obtained by our introduced approaches as well as network violations, in normal operating conditions (Energy) and when either raise or lower reserve market gets called (Reserve). We also report the number of times that the available FCAS capacity was less than the bid submitted to the market. As reported, PWA-ARCO obtains 2.6% less benefit compared to the perfect yet unachievable case. Notice that our approach requires 5 times less computation and communication than Perfect. Plus, unlike Perfect, we do not have all the future information. It might seem counter intuitive that Deterministic has obtained benefit close to Perfect. However, here we have not applied any penalty for reserve violation and as reported in Table I, in 372 scenarios out of our total 1440 scenarios, Deterministic was not able to honour the accepted reserve bids. In addition, as reported in Table I, Deterministic violates the safe voltage limits in 341 scenarios. Fig. 5 provides a visualisation on maximum voltages across the network using Deterministic and PWA-ARCO.

TABLE I: Total Benefit, Network / Reserve Bid Violations

Approach	Benefit ($)	Rel. to Per. (%)	#Violation		
			Network		Reserve Bids
			Energy	Reserve	
Deterministic	3502.5	-2.0	92	249	372
Perfect	3574.9	–	0	0	0
AARCO	2746.2	-23	0	0	0
PWA-ARCO	3481.5	-2.6	0	0	0

While AARCO could keep the voltages within the safe limits and honour its FCAS bids, it adds a significant cost to the bidding problem. The reason is that, compared to a piecewise function, compensating uncertainty with affine functions will lead to over-conservative results. We further study this in Section VI-B. Also, remember that we relaxed binary variables (8a)–(8b) which can potentially make PWA-ARCO more conservative. However, notice that the perfect case does not count on PWA-ARCO local control and solve optimisation every minute to find the optimum solution. Therefore, it serves as the best upper bound to the benefit we could obtain. In our case, our PWA-ARCO could get as close as 2.6% to this perfect yet unachievable solution. Although this gap might vary in different cases, it still provides a good insight on how much room there is for improvement.

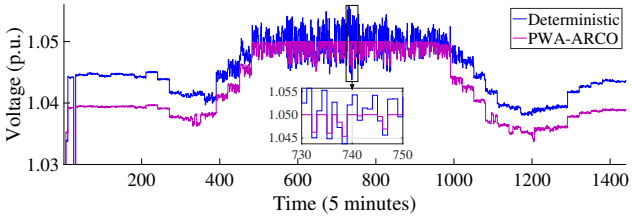

Fig. 5: Maximum voltages across the grid

Fig. 6 breaks the total benefit obtained by each approach down to the energy and reserve market components. As can be seen in Fig. 6, the majority of the total benefit for all cases is made through the FCAS market. The reason is that unlike in reserve markets, consumers have a cost in the energy market, e.g., to meet their background load / charge their batteries. It is worth mentioning that the revenue in reserve market could be lost in the deterministic approach due to the likelihood of failing to fulfil the reserve bids.

[5]We found these values to provide a good convergence. Investigating on ADMM penalty parameters is out of the scope of this study.

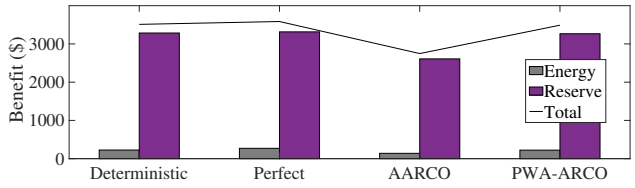

Fig. 6: Energy and FCAS benefits breakdown

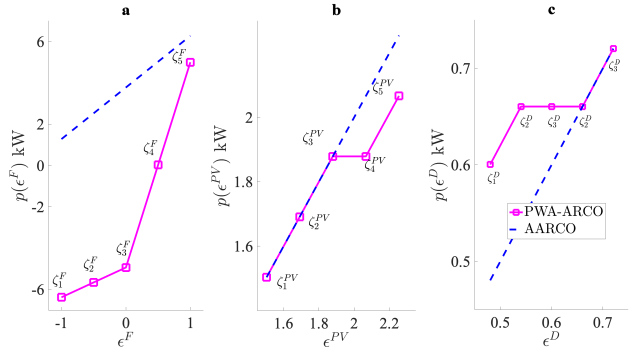

Fig. 7: An example of CPP function: PWA-ARCO vs. AARCO

B. PWA-ARCO vs. AARCO

As we mentioned before, instead of optimising for affine functions as in AARCO, we optimise for piecewise affine functions which enables a better response to uncertainty. To compare the piecewise affine policy with just affine, we plot the optimised functions for CPP obtained by AARCO and PWA-ARCO approaches for a random consumer at 1:30 PM in Fig. 7. Subfigures **a**, **b**, **c** are respectively associated with the first (reserve activation), second (PV) and third (demand) terms in (13a). As plotted, AARCO compensates for uncertainty with the same slope (i.e. A in $A\epsilon + b$), yet PWA-ARCO responds differently for different part of the uncertainty set to obtain higher benefits. For instance, regarding Fig. 7. **a**, notice that in AARCO the reserve capacity is limited to the same slope (2.5 kW raise, 2.5kW lower). On the other hand, PWA-ARCO is able to allocate the flexibility differently within the uncertainty set, leading to 2kW lower (i.e. $p(\zeta_3^F) - p(\zeta_1^F)$) and 10kW raise (i.e. $p(\zeta_5^F) - p(\zeta_3^F)$) offers.

C. ADMM envelope vs. Fixed Envelopes

To study the effectiveness of accounting for consumer preferences in calculation of operating envelopes, here, we study 3 other approaches in which the DSO obtains the operating envelope and allocate them to consumers without our ADMM negotiation tool. We study the following envelopes:

Equal-Width Envelopes: In which the DSO allocates equal-width envelopes to all consumers owning DER.

Max Export Envelopes [2]: In which the DSO allocates envelopes of different widths to consumers at different nodes such that the absolute network throughput is maximised.

Fair Envelopes: We use the objective function (1) in [2] to obtain fair operating envelopes for each consumer.

Consumers use these envelopes as hard constraints to obtain their bids. To have a fair comparison, we used PWA-ARCO in consumer subproblems for all approaches. Table II reports the benefit obtained via the equal-width, max export and fair envelopes vs. our proposed distributed envelope. Since our approach continuously negotiates and obtains the operating envelopes, it attains the best results in total. Notice that in line with [3], the equal-width, max-export and fair envelopes are obtained assuming that the DSO has full observability of consumers' DER. However, in reality, such observability may not be available to the DSO, resulting in less representative envelopes further reducing consumers' benefit.

TABLE II: Dynamic Operating Envelopes vs. Proposed

Envelopes	Benefits ($)			Rel. to Proposed %
	Energy	Reserve	Total	
Equal-Width	226.6	2956.3	3182.9	-8.5
Max-Export	208.5	3072.8	3281.3	-5.6
Fair	212.7	3098.6	3311.3	-4.8
Proposed	214.8	3266.8	3481.6	–

D. Longer Term Benefit

To evaluate the longer-term performance of PWA-ARCO, we have run our approach for 3 months, from 1st January to 31st March 2021. Table III reports the benefit obtained by Deterministic, PWA-ARCO as well as a case in which consumers use PWA-ARCO to participate only in the energy market. To ensure that PWA-ARCO covers the real-data variations over three months, we have chosen the uncertainty sets to cover $\pm 35\%$ deviation about the forecasts. Notice that smaller uncertainty sets can be chosen based on the month or the day of the week. However, here we keep the uncertainty set constant for all days. PWA-ARCO could both honour the reserve bids and keep the voltages within the accepted range. We have also run PWA-ARCO with a constant $\pm 20\%$ uncertainty coverage, yet in few operating points, such an uncertainty set was not big enough to cover all scenarios. In our test case, with $\pm 20\%$ uncertainty coverage, we obtained 5% higher benefit (70,644.5 compared to 70,314.3), yet in 105 operating points out of the total of 78k, the network was violated.

TABLE III: Three-Month Benefit

Approach	Total Benefit ($)	Violation	
		Network	Reserve
Energy	-3,344.0	–	–
Deterministic	70,396.6	3,883	9,261
PWA-ARCO	70,314.3	–	–

In our experiment, Deterministic violated both the network constraints (in 3.9k out of the total 78k operating points) and reserve bids (in 9.3k out of 155.5k reserve bids). Notice that in the NEM, not honouring reserve bids can exclude consumers from these highly beneficial markets. We also report the benefit that could be obtained in the energy market alone. As can be seen, if consumers cannot bid into reserve markets their benefit would shrink significantly.

E. Convergence and Problem Size

Table IV reports the number of optimisations that needs to be solved within every ADMM iteration in the consumer

and the DSO subproblem, the number of variables[6] at a single optimisation, as well as the total number of ADMM iterations.

TABLE IV: Problem Size and Convergence

Approach	No. Iter	No. of Problems × Var. in a Problem	
		Consumer	DSO
Deterministic	66	207 × 4.5k	564 × 482
Perfect	73	207 × 27.4k	2880 × 482
AARCO	33	207 × 5.2k	564 × 482
PWA-ARCO	10	207 × 12.9k	564 × 482

Notice that our ADMM approach enables us to use parallel computing in both consumer and network subproblems. For example in the consumer subproblem, every consumer solves its optimisation problem at the same time on their EMS. Thus, the consumer sub-problem is solved fully in parallel in a real-world setting. Similarly, the OPFs for every time step can be solved in parallel. Thus, the solve time of each subproblem at every iteration is equal to the longest time taken by a single optimisation within the subproblem. In our implementation, we solve the problem of every consumer and OPF in a loop one after another and use the longest time at every loop as the time taken by the parallel real-world case. We report the time taken by our sequential as well as the equivalent parallel run time in Table V. The reported run times are obtained using a laptop computer with a 2.50 GHz Intel(R) Core(TM) i7 and 8 GB of memory.

TABLE V: Computation Time

Approach	Parallel (Sequential) Implementation		
	Single Iteration		Total
	Consumer	Network	
Deterministic	0.08s (7.6s)	0.14s(23.2s)	14.1s (33.9m)
Perfect	0.24s (45.3s)	0.14s(118.2s)	27.8s (198.9m)
AARCO	0.09s (9.3s)	0.14s(22.6s)	7.7s (17.6m)
PWA-ARCO	0.18s (29.4s)	0.14s(24.4s)	3.1s (9.0m)

VII. Conclusion

We obtained network-secure operating envelopes which facilitate market participation of consumers. To ensure that consumers can commit to their envelopes, account for bid violation penalties, and generate reliable reserve bids, we introduce a novel PWA-ARCO consumer bidding. Both the envelopes and parameters of our piecewise affine controllers are obtained during the negotiation of the ADMM approach. This tunes our local controllers based on global measurements and helps improve the results. In live operation, PV power and demand are continually measured and together with any FCAS activation signals are fed into our PWA-ARCO controllers upon which proper recourse actions are taken. This is a really valuable feature of our approach as it can reduce the need for frequent negotiations making our approach more functional in practice. To increase network throughput, we enable consumers to grant the grid with their reactive power support. Our results demonstrate that our PWA-ARCO approach serves its purpose whilst providing an excellent compromise between computational cost and the solution quality.

[6]The number of constraints in our optimisations are approximately equal to the number of variables.

Appendix A
Proof of Remark 1

Here, we aim to prove that if the network is radial and operates within the voltage stable region, then solving two OPFs, associated with the extremes, captures the worst conditions.

Proof: Let us begin by linearising the power flow equations about an operating point, which in our case can be considered the forecasted values. Let $\tilde{V}_i$ denote the voltage magnitude at node i at an initial operating point, and let Δp_i denote the deviation of real power injection at node i from its value at that operating point. Also, let K^p be the sensitivity matrix, where K^p_{ij} denotes the voltage magnitude sensitivity of node i to real power injection at node j. Assuming zero reactive power changes[7], we write the following relation:

$$V_i = \tilde{V}_i + \sum_{j=1}^{N} K^p_{ij} \Delta p_j, \tag{A.1}$$

where V_i denotes the voltage magnitude at node i; and N is the number of nodes in the network. We wish to obtain the conditions under which the voltage magnitude at all the nodes stay within a desired bound for all realisations of the nodal real power deviations within a box uncertainty set. More formally:

$$\underline{V_i} \leq \tilde{V}_i + \sum_{j=1}^{N} K^p_{ij} \Delta p_j \leq \overline{V_i} \quad \forall \Delta p_j \in \mathcal{U}_j, i \in \{1, ..., N\}. \tag{A.2}$$

where $\mathcal{U}_j = [\underline{\Delta p_j}, \overline{\Delta p_j}]$, and $\underline{V_i}$ and $\overline{V_i}$ denote the acceptable lower and upper voltage limits at note i, respectively. Using *max* and *min* protection functions, we rewrite (A.2) as:

$$\tilde{V}_i + \max_{\Delta p_j \in \mathcal{U}_j} \{\sum_{j=1}^{N} K^p_{ij} \Delta p_j\} \leq \overline{V_i} \tag{A.3a}$$

$$\tilde{V}_i + \min_{\Delta p_j \in \mathcal{U}_j} \{\sum_{j=1}^{N} K^p_{ij} \Delta p_j\} \geq \underline{V_i} \tag{A.3b}$$

We now assume that K^p is a non-negative matrix, i.e., we assume that all its elements are equal to or greater than zero (in one moment we will discuss the requirements for this assumption to hold). Thus, we rewrite (A.3) as the following:

$$\tilde{V}_i + \sum_{j=1}^{N} K^p_{ij} \overline{\Delta p_j} \leq \overline{V_i} \tag{A.4a}$$

$$\tilde{V}_i + \sum_{j=1}^{N} K^p_{ij} \underline{\Delta p_j} \geq \underline{V_i}. \tag{A.4b}$$

Hence under the assumption that K^p is a non-negative matrix, to make sure that all the combinations of real power deviations will not lead to voltage violations, we only need to check for the two scenarios where all the injections are simultaneously on either side of the uncertainty set $\mathcal{U}_j$.

We now check the requirements for our assumption. Based on the voltage stability definition [32], Chapter 14, if the

[7]The same proof holds when the reactive power is included. Because reactive power is controllable here and is only meant to improve the voltages.

sensitivity of voltage magnitude to real power injection becomes negative the system is voltage unstable. Hence our first requirement is working in voltage stable operating region, which is a fair assumption in a power system operation setting.

Notice that in the stability argument, we are only looking at the diagonal elements in the sensitivity matrix. For the rest of the elements in the sensitivity matrix, we look at the physical structure of nodal connections. For a radial networks, neglecting the higher order real and reactive power loss terms, we can write the power flow equations as in (A.5). This approximation introduces a small relative error of at most 0.25% if there is a 5% deviation in voltage magnitude.

$$V_i = V_0 + \sum_j R_{ij} p_j, \tag{A.5}$$

where V_0 denotes the voltage magnitude at the slack node, and R_{ij} denotes the resistance of the direct path between nodes i and j. Comparing (A.5) and (A.1), we can see that R_{ij} is an approximation of K^p_{ij}, and indeed, is equal to it when we are linearising the power flow equation at the no-load condition. Notice that since the elements in R are the summation of positive resistance of the lines on the path between the nodes, they all have positive values. Hence, in radial system, the assumption of non-negative sensitivity is valid. A similar proof can be written for the thermal limits of the network.

REFERENCES

[1] L. Blackhall, "On the calculation and use of dynamic operating envelopes," ARENA Knowledge Sharing Report, Tech. Rep., 2020, https://arena.gov.au/assets/2020/09/on-the-calculation-and-use-of-dynamic-operating-envelopes.pdf.

[2] K. Petrou, M. Z. Liu, A. T. Procopiou, L. F. Ochoa, J. Theunissen, and J. Harding, "Operating envelopes for prosumers in LV networks: A weighted proportional fairness approach," in *2020 IEEE PES Innovative Smart Grid Technologies Europe (ISGT-Europe)*, 2020.

[3] K. Petrou, A. Procopiou, L. Gutierrez-Lagos, M. Liu, L. Ochoa, T. Langstaff, and J. Theunissen, "Ensuring distribution network integrity using dynamic operating limits for prosumers," *IEEE Transactions on Smart Grid*, 2021.

[4] A. Ben-Tal, A. Goryashko, E. Guslitzer, and A. Nemirovski, "Adjustable robust solutions of uncertain linear programs," *Mathematical Programming*, vol. 99, no. 2, pp. 351–376, 2004.

[5] K. Oikonomou, M. Parvania, and R. Khatami, "Coordinated deliverable energy flexibility and regulation capacity of distribution networks," *International Journal of Electrical Power & Energy Systems*, vol. 123, p. 106219, 2020.

[6] A. Safdarian, M. Fotuhi-Firuzabad, and M. Lehtonen, "Optimal residential load management in smart grids: A decentralized framework," *IEEE Transactions on Smart Grid*, vol. 7, no. 4, pp. 1836–1845, 2016.

[7] Y. Wang, L. Wu, and S. Wang, "A fully-decentralized consensus-based ADMM approach for DC-OPF with demand response," *IEEE Transactions on Smart Grid*, vol. 8, no. 6, pp. 2637–2647, 2017.

[8] S. Mhanna, G. Verbič, and A. C. Chapman, "Adaptive ADMM for distributed AC optimal power flow," *IEEE Transactions on Power Systems*, vol. 34, no. 3, pp. 2025–2035, 2019.

[9] P. Scott, D. Gordon, E. Franklin, L. Jones, and S. Thiébaux, "Network-aware coordination of residential distributed energy resources," *IEEE Transactions on Smart Grid*, 2019.

[10] T. R. Ricciardi, K. Petrou, J. F. Franco, and L. F. Ochoa, "Defining customer export limits in PV-rich low voltage networks," *IEEE Transactions on Power Systems*, vol. 34, no. 1, pp. 87–97, 2018.

[11] A. Attarha, P. Scott, and S. Thiébaux, "Network-aware co-optimisation of residential DER in energy and FCAS markets," *PSCC 2020*, 2020.

[12] A. Attarha, P. Scott, J. Iria, and S. Thiébaux, "Network-secure and price-elastic aggregator bidding in energy and reserve markets," *IEEE Transactions on Smart Grid*, vol. 12, no. 3, pp. 2284–2294, 2021.

[13] Y. Wang, H. Liang, and V. Dinavahi, "Two-stage stochastic demand response in smart grid considering random appliance usage patterns," *IET Generation, Transmission & Distribution*, vol. 12, no. 18, pp. 4163–4171, 2018.

[14] F. U. Nazir, B. C. Pal, and R. A. Jabr, "Distributed solution of stochastic volt/var control in radial networks," *IEEE Transactions on Smart Grid*, 2020.

[15] J. Wang, M. Shahidehpour, and Z. Li, "Security-constrained unit commitment with volatile wind power generation," *IEEE Transactions on Power Systems*, vol. 23, no. 3, pp. 1319–1327, 2008.

[16] C. He, L. Wu, T. Liu, and M. Shahidehpour, "Robust co-optimization scheduling of electricity and natural gas systems via ADMM," *IEEE Transactions on Sustainable Energy*, vol. 8, no. 2, pp. 658–670, 2017.

[17] R. Jiang, J. Wang, and Y. Guan, "Robust unit commitment with wind power and pumped storage hydro," *IEEE Transactions on Power Systems*, vol. 27, no. 2, pp. 800–810, 2011.

[18] J. S. Giraldo, J. A. Castrillon, J. C. López, M. J. Rider, and C. A. Castro, "Microgrids energy management using robust convex programming," *IEEE Transactions on Smart Grid*, 2018.

[19] R.-S. Liu and Y.-F. Hsu, "A scalable and robust approach to demand side management for smart grids with uncertain renewable power generation and bi-directional energy trading," *International Journal of Electrical Power & Energy Systems*, vol. 97, pp. 396–407, 2018.

[20] R. Louca and E. Bitar, "Robust AC optimal power flow," *IEEE Transactions on Power Systems*, vol. 34, no. 3, pp. 1669–1681, 2018.

[21] R. A. Jabr, "Robust Volt/VAr control with photovoltaics," *IEEE Transactions on Power Systems*, vol. 34, no. 3, pp. 2401–2408, 2019.

[22] A. Zare, C. Chung, J. Zhan, and S. O. Faried, "A distributionally robust chance-constrained milp model for multistage distribution system planning with uncertain renewables and loads," *IEEE Transactions on Power Systems*, vol. 33, no. 5, pp. 5248–5262, 2018.

[23] R. Mieth and Y. Dvorkin, "Data-driven distributionally robust optimal power flow for distribution systems," *IEEE Control Systems Letters*, vol. 2, no. 3, pp. 363–368, 2018.

[24] B. K. Poolla, A. R. Hota, S. Bolognani, D. S. Callaway, and A. Cherukuri, "Wasserstein distributionally robust look-ahead economic dispatch," *IEEE Transactions on Power Systems*, vol. 36, no. 3, pp. 2010–2022, 2020.

[25] A. Attarha, P. Scott, and S. Thiébaux, "Affinely adjustable robust ADMM for residential DER coordination in distribution networks," *IEEE Transactions on Smart Grid*, 2019.

[26] H. Gao, J. Liu, L. Wang, and Z. Wei, "Decentralized energy management for networked microgrids in future distribution systems," *IEEE Transactions on Power Systems*, vol. 33, no. 4, pp. 3599–3610, 2017.

[27] M. Mohiti, H. Monsef, A. Anvari-Moghaddam, J. Guerrero, and H. Lesani, "A decentralized robust model for optimal operation of distribution companies with private microgrids," *International Journal of Electrical Power & Energy Systems*, vol. 106, pp. 105–123, 2019.

[28] X. Chen, E. Dall'Anese, C. Zhao, and N. Li, "Aggregate power flexibility in unbalanced distribution systems," *IEEE Transactions on Smart Grid*, vol. 11, no. 1, pp. 258–269, 2019.

[29] F. Capitanescu, "TSO–DSO interaction: Active distribution network power chart for TSO ancillary services provision," *Electric Power Systems Research*, vol. 163, pp. 226–230, 2018.

[30] S. Boyd, N. Parikh, E. Chu, B. Peleato, J. Eckstein *et al.*, "Distributed optimization and statistical learning via the alternating direction method of multipliers," *Foundations and Trends® in Machine learning*, vol. 3, no. 1, pp. 1–122, 2011.

[31] I. Dunning, J. Huchette, and M. Lubin, "Jump: A modeling language for mathematical optimization," *SIAM Review*, vol. 59, no. 2, pp. 295–320, 2017.

[32] P. Kundur, "Power system stability," *Power system stability and control*, pp. 7–1, 2007.