Privacy-Preserving Distributed Market Mechanism for Active Distribution Networks

Matthias Franke*, Ognjen Stanojev*, Lesia Mitridati†, Gabriela Hug*
*ETH - Power Systems Laboratory, ETH Zürich, Zürich, Switzerland
† Center for Electric Power and Energy, Technical University of Denmark, Kgs. Lyngby, Denmark
{mfranke, ognjens, ghug}@ethz.ch, lemitri@dtu.dk

***Abstract*—Amidst the worldwide efforts to decarbonize power networks, Local Electricity Markets (LEMs) in distribution networks are gaining importance due to the increased adoption of renewable energy sources and prosumers. Considering that LEMs involve data exchange among independent entities, privacy and cybersecurity are some of the main practical challenges in LEM design. This paper proposes a secure market protocol using innovations from distributed optimization and Secure MultiParty Computation (SMPC). The considered LEM is formulated as an uncertainty-aware joint market for energy and reserves with affine balancing policies. To achieve scalability and enable the use of SMPC, market clearing is solved using the Consensus ADMM algorithm. Subsequently, the data exchange among participants via ADMM iterations is protected using the Shamir secret-sharing scheme to ensure privacy. The market protocol is further reinforced by a secure and verifiable settlement process that uses SMPC and ElGamal commitments to verify market quantities and by a secure recovery scheme for missing network measurements. Finally, the feasibility and performance of the proposed LEM are evaluated on a 15-bus test network.**

***Index Terms*—Consensus ADMM, Local Electricity Markets, Cyber Security, Secure Multiparty Computation**

I. Introduction

With the goal to mitigate climate change, an increasing effort is made to decarbonize power networks. This effort primarily involves the increased adoption of Distributed Energy Resources (DERs) in distribution grids and empowering households to become *prosumers* that can contribute to decarbonization [1]. Despite these advancements, current electricity markets still operate in a hierarchical, top-down manner, without adapting their market structures to the emergence of DERs [2]. Furthermore, technical challenges arise in the operation of distribution networks due to high bi-directional power flows and voltage fluctuations [3], [4]. These challenges suggest the development of Local Electricity Markets (LEMs) in which small-scale producers and consumers can transact electricity in a decentralized fashion while respecting the physical limits of the distribution network. Nevertheless, such a collaborative principle involves data exchange among independent entities, therefore rising privacy and cybersecurity concerns. We propose an uncertainty-aware LEM that is operated in a secure and distributed manner using innovations from optimization and theoretical cryptography.

There has been a wide range of research into innovating LEMs that promise to safely and efficiently operate future networks and markets with DERs [1]. The three main approaches for such LEMs are [2]: (i) pure Peer-to-Peer markets, whose fully decentralized approach offers customers high flexibility at the cost of lacking convergence and safety guarantees, (ii) community markets with an operator efficiently managing trading both locally and with the upper-level grid, however, resulting in computational scaling issues, and (iii) hybrid methods that use a bilevel approach to leverage the advantages of the former two. Of key interest for this paper are the community markets, which are proven to successfully operate LEMs with DERs in the existing literature [5], [6]. However, these markets have encountered challenges due to the use of oversimplified network models [5], the absence of DER uncertainty management [2], and the lack of effective coordination with the upper-level grid [6]. Additionally, most works in this field focus on theoretical explorations of such markets, and there is thus a lack of research into how to provide cyber security and resilience to measurement failure for the actual implementation of LEMs.

Efforts to overcome these gaps have been made, although they have not been consolidated yet. For instance, improved models of distribution networks together with techniques to capture the stochastic nature of DERs are available [7], such as the LinDistFlow formulation [8], [9] that achieves high accuracy through the utilization of linear and convex constraints. This formulation overcomes shortcomings of the DC power flow formulations for distribution networks [8] and is further highly suitable for adding chance constraints to capture the uncertainty of DERs within a market model [6]. Of particular interest to this paper is the usage of chance constraints to concurrently optimize day-ahead energy and reserve needs [5].

To overcome computational complexity and scalability issues, the usage of distributed optimization for the solution of LEM formulations is suitable and further allows the integration of the relevant cryptographic techniques. The work in [10] decomposes the collaborative optimization problem using Lagrangian relaxation and solves the master problem via Secure Multi-Party Computation (SMPC) with secret sharing protocols. However, it mainly serves as a proof of concept, with intentionally simple network and market models, and the authors themselves highlight the Lagrangian Relaxation's inability to provide convergence guarantees. Such guarantees can be provided using the Consensus version of the Alternating Direction Method of Multipliers (ADMM) [11], which has already been used with the LinDistFlow grid model [8], [12].

The main cryptographic technique employed within this study is SMPC, which belongs to a family of cryptographic protocols that can execute arbitrary multiparty computations with information-theoretic security in a threshold model [13]. This characteristic renders SMPC well-suited for use in LEM design. Notably, it sidesteps the challenges that previous privacy-preserving optimization methods encountered, such as poor computational performance, degraded convergence guarantees, and weaker security notions like differential privacy [14], [15]. Furthermore, the applicability of SMPC extends to reinforcing the security of LEMs protecting against network sensor failures due to potential interference of adversaries [16].

This paper proposes a fully privacy-preserving LEM framework based on distributed optimization and SMPC with secret sharing protocols. Building upon prior research [6], we develop an uncertainty-aware joint market for energy and reserves using the chance-constrained LinDistFlow formulation. The model is further extended by introducing batteries as local energy storage [17] and by capturing the tariff-switching behavior of the substation [18]. To overcome the computational limitations encountered in [10], we propose a distributed and privacy-preserving market clearing mechanism using an SMPC-secured Consensus ADMM algorithm, as described in Sec. II-D and Sec. III-D. The secure version of the LEM requires mechanisms to preserve the security of the system between SMPC sessions, for which we use standard techniques called commitment schemes and Zero-Knowledge Proofs that can be applied without impinging on the overall privacy preservation [19]. In particular, we leverage SMPC and ElGamal commitments [20] to provide a double verification scheme and a measurement recovery scheme (Sec. III-E) to provide a secure and verifiable settlement process (Sec. III-F). This setup allows the proposed market mechanism to achieve similar solutions as a central insecure solver but with added privacy preservation, as shown by the results in Sec. IV.

II. Local Electricity Market Design

A. Preliminaries

This paper studies a LEM in an active distribution grid. The distribution network is represented as an undirected and connected tree graph $\mathcal{G}(\mathcal{N}, \mathcal{E})$, where $\mathcal{N} = \{0, 1, \dots, N\}$ is the set of nodes in the graph and $\mathcal{E} \subset \mathcal{N} \times \mathcal{N}$ is the set of N edges. The substation node is indexed by 0 and represents the interface between the distribution network and the upper-level grid. Furthermore, we use $\mathcal{N}^+ = \mathcal{N} \setminus \{0\}$ to denote the set of non-substation nodes. For a node $n \in \mathcal{N}$, its ancestor node is denoted by A_n, while the set of its children nodes is denoted by $\mathcal{C}_n$. The notation related to the physical quantities of the LinDistFlow network model is introduced in the following.

For each node $n \in \mathcal{N}^+$, let u_n denote the squared voltage magnitude at the node, and $u_n^{\max}$ and $u_n^{\min}$ the corresponding maximum and minimum voltage limits. The substation node is assumed to have a fixed predefined voltage magnitude u_0. The active and reactive line flows to a node $n \in \mathcal{N}^+$ from its ancestor node A_n are denoted by f_n^P and f_n^Q, respectively. Furthermore, each such line is characterized by a resistance r_n, a reactance x_n, and a maximum apparent power limit S_n.

The participants in the considered LEM include: (i) consumers at all non-substation nodes $i \in \mathcal{N}^+$, who each have known and inflexible active d_i^P and reactive d_i^Q demand profiles, (ii) prosumers $r \in \mathcal{R} \subseteq \mathcal{N}^+$ who own DERs with a forecasted active power production h_r^f, (iii) batteries $m \in \mathcal{M} \subseteq \mathcal{N}^+$, characterized by their State of Charge (SoC) B_m, and the corresponding maximum $B_m^{\max}$ and minimum $B_m^{\min}$ SOC limits, and (iv) the substation node, which is centrally operated to trade energy and procure reserves with the wholesale markets. Batteries and prosumers constitute flexible generators $v \in \mathcal{V} = \mathcal{R} \cup \mathcal{M}$ with adjustable active power generation g_v^P between maximum $P_v^{\max}$ and minimum $P_v^{\min}$ limits. The cost of active power adjustment is characterized by a quadratic c_v^q and a linear c_v^l cost coefficient. The substation node is the "infinite bus", with its generation modeled as the difference between the active l^P and reactive l^Q inflow and active s^P and reactive s^Q outflow to capture a tariff-switching scheme without the need for binary variables [18], i.e., the system operator charges inflow to the LEM by the forecasted wholesale price plus a flat usage tariff Φ^+, and pays for outflow at the wholesale price minus the tariff Φ^-.

B. Uncertainty Modeling

The DER generation of a prosumer $r \in \mathcal{R}$ introduces uncertainty in the LEM and is thus modeled by

$$h_r = h_r^f + \omega_r, \tag{1}$$

where h_r is the stochastic DER generation and ω_r is the random forecast error. The forecast error is modeled as an independent Gaussian random variable with zero mean $\mathbb{E}[\omega_r] = 0$ and variance $\mathrm{Var}[\omega_r] = \sigma_r$, known only to the prosumer $r \in \mathcal{R}$ itself. The total DER forecast error in the system can thus be calculated as $\Delta = \sum_{r \in \mathcal{R}} \omega_r$, resulting in a zero-mean multivariate distribution with $\Sigma = \mathrm{Var}[\Delta] = \mathrm{diag}(\{\sigma_r\}_{r \in \mathcal{R}})$.

The realized total forecast error creates a power imbalance in the system that requires a global response from the flexible assets. To this end, we equip each flexible generator $v \in \mathcal{V}$ with a linear reserve policy [5], characterized by a participation factor α_v, such that its total generation is given by

$$\tilde{g}_v^P = g_v^P - \alpha_v \Delta. \tag{2}$$

Given that $\sum_{v \in \mathcal{V}} \alpha_v = 1$, the flexible generators completely balance the observed overall active power deviation.

C. Market Clearing Formulation

In this section, we present the chance-constrained market formulation. The variable set includes the previously defined network and adjustable generation quantities and is defined by

$$\begin{aligned}\Xi^{\mathbb{P}} = &\{u_{i,t}, f_{i,t}^P, f_{i,t}^Q\}_{i \in \mathcal{N}^+, t \in \mathcal{T}} \cup \{g_{v,t}^P, \alpha_{v,t}\}_{v \in \mathcal{V}, t \in \mathcal{T}} \\ &\cup \{l_t^P, s_t^P, l_t^Q, s_t^Q\}_{t \in \mathcal{T}} \cup \{B_{m,t}\}_{m \in \mathcal{M}, t \in \mathcal{T}},\end{aligned} \tag{3}$$

where $\mathcal{T}$ is the set of considered time steps. The aim of the proposed local market is to minimize the expected cost of

generation and energy procurement from the upper-level grid while at the same time respecting the network and generation constraints for all time steps $t \in \mathcal{T}$:

$$\min_{\Xi^{\mathbb{P}}} \mathbb{E}\Big[\sum_{t\in\mathcal{T}}\sum_{v\in\mathcal{V}}\Big(c_v^q(\tilde{g}_{v,t}^P)^2 + c_r^l \tilde{g}_{v,t}^P\Big) + l_t^P\Phi_t^+ - s_t^P\Phi_t^-\Big] \quad (4a)$$

$$\text{s.t.} \quad l_t^P - s_t^P = \sum_{j\in\mathcal{C}_0} f_{j,t}^P \quad (4b)$$

$$l_t^Q - s_t^Q = \sum_{j\in\mathcal{C}_0} f_{j,t}^Q \quad (4c)$$

$$l_t^P \geq 0, s_t^P \geq 0, l_t^Q \geq 0, s_t^Q \geq 0 \quad (4d)$$

$$f_{n,t}^P + (\tilde{g}_{n,t}^P - d_{n,t}^P + h_{n,t}) = \sum_{j\in\mathcal{C}_n} f_{j,t}^P, \qquad \forall n \in \mathcal{N}^+ \quad (4e)$$

$$f_{n,t}^Q - d_{n,t}^Q = \sum_{j\in\mathcal{C}_n} f_{j,t}^Q, \qquad \forall n \in \mathcal{N}^+ \quad (4f)$$

$$u_{n,t} = u_{A_n,t} - 2(r_n f_{n,t}^P + x_n f_{n,t}^Q), \qquad \forall n \in \mathcal{N}^+ \quad (4g)$$

$$B_{m,t} = B_{m,t-1} - \tilde{g}_{m,t}^P, \qquad \forall m \in \mathcal{M} \quad (4h)$$

$$\mathbb{P}[u_{n,t} \leq u_n^{\max}] \geq 1 - \epsilon_u, \qquad \forall n \in \mathcal{N}^+ \quad (4i)$$

$$\mathbb{P}[u_n^{\min} \leq u_{n,t}] \geq 1 - \epsilon_u, \qquad \forall n \in \mathcal{N}^+ \quad (4j)$$

$$\mathbb{P}[a_n^1 f_{n,t}^P + a_n^2 f_{n,t}^Q + a_n^3 S_n \leq 0] \geq 1 - \epsilon_f, \quad \forall n \in \mathcal{N}^+ \quad (4k)$$

$$\mathbb{P}[\tilde{g}_{v,t}^P \leq P_v^{\max}] \geq 1 - \epsilon_g, \qquad \forall v \in \mathcal{V} \quad (4l)$$

$$\mathbb{P}[P_v^{\min} \leq \tilde{g}_{v,t}^P] \geq 1 - \epsilon_g, \qquad \forall v \in \mathcal{V} \quad (4m)$$

$$\mathbb{P}[B_{m,t} \leq B_m^{\max}] \geq 1 - \epsilon_b, \qquad \forall m \in \mathcal{M} \quad (4n)$$

$$\mathbb{P}[B_m^{\min} \leq B_{m,t}] \geq 1 - \epsilon_b, \qquad \forall m \in \mathcal{M} \quad (4o)$$

$$\sum_{v\in\mathcal{V}} \alpha_{v,t} = 1, \quad 0 \leq \alpha_{v,t} \leq 1, \qquad \forall v \in \mathcal{V}. \quad (4p)$$

Constraints related to inflow/outflow at the substation node are given in (4b)-(4d). The LinDistFlow network model is established in (4e)-(4g), and the battery SoC model[1] in (4h). The chance constraints on bus voltages are enforced in (4i)-(4j), power generation limits in (4l)-(4m), battery SoC in (4n)-(4o), and the dodecagon linear approximations (defined by coefficients $a_n^1, a_n^2, a_n^3, \forall n \in \mathcal{N}^+$) of the line flow constraints in (4k). The LEM operator specifies an error term for each type of CC, representing the maximum acceptable percentage of a constraint violation, namely, ϵ_g for generation limits, ϵ_b for battery limits, ϵ_u for voltage limits and ϵ_f for flow limits. Finally, the bounds on reserve coefficients are given in (4p).

D. Distributed Solution Method

To obtain a tractable form of the optimization problem in (4), each individual linear chance constraint can be reformulated as a second-order cone constraint. This process is omitted for brevity, and we refer the reader to [6] for more details. Instead, we here focus on decomposing the centralized optimization problem into a distributed form using the scaled formulation of Consensus ADMM [11], as proposed in [12].

[1]For the sake of simplicity, but without loss of generality, we assume perfect charging and discharging efficiencies.

Considering that the objective function in (4a) is separable, let us introduce a cost function $f_n(X_n)$ related to each node $n \in \mathcal{N}$, where X_n denotes the vector of all variables related to node n. A subset of these variables, collected in $X_{\mathbb{C}_n} = (u_n, f_n^P, f_n^Q, \alpha_n)$ and called coupling variables, appear in the constraints of other nodes. The coupling variables have a global copy $Z_{\mathbb{C}_n}$ in the Consensus ADMM algorithm that ensures their system-wide convergence. The full algorithm at iteration k is given by

$$X_{\mathbb{C}_n}^{k+1} = \arg\min f_n(X_n) + \frac{\rho_0}{2}||X_{\mathbb{C}_n}^k - Z_{\mathbb{C}_n}^k + U_{\mathbb{C}_n}^k||_2^2, \quad (5a)$$

$$Z_{\mathbb{C}_n}^{k+1} = \mathbb{AVG}(X^{k+1}), \quad (5b)$$

$$U_{\mathbb{C}_n}^{k+1} = U_{\mathbb{C}_n}^k + X_{\mathbb{C}_n}^{k+1} - Z_{\mathbb{C}_n}^{k+1}, \quad (5c)$$

where ρ_0 is a positive scaling factor and $U_{\mathbb{C}_n}$ are the Lagrange multipliers. In (5b), an arithmetic mean of all relevant local values is computed to find the global values of the coupled variables. The employed convergence metrics are the ℓ_2-norms of the primal and dual residuals of the local problems and the total power surplus in the system relative to the total demand.

E. Financial Settlement

Solving the market problem (4) yields nodal λ_n and flexibility π prices, calculated as dual variables of constraints (4e) and (4p), respectively. Using these prices, the financial payoff for a market participant n from the daily market operation is

$$\mathcal{B}_n = \sum_{t\in\mathcal{T}}\Big[\lambda_{n,t}^P(g_{n,t}^P + h_{n,t}^f - d_{n,t}^P) + \pi_t(\alpha_{n,t} - \frac{\sigma_{n,t}}{\sum_{r\in\mathcal{R}}\sigma_{r,t}})\Big] \quad (6)$$

As can be seen above, the participants are charged at the flexibility price π for their share of the overall uncertainty, captured by the standard deviation of their forecast error, as well as a fixed tariff for the power flow via the substation [18]. Their profit stems from the provided flexibility and the injected DER controllable and forecasted active power. The final balance $\mathcal{F}_n$ of a party n, however, also includes the results from the two-price imbalance regulation $\mathcal{I}_n$ performed by the DSO [21] on any remaining deviations:

$$\mathcal{F}_n = \mathcal{B}_n + \mathcal{I}_n. \quad (7)$$

III. Secure Local Market Protocol

A. Secure Multiparty Computation Preliminaries

The technique used to ensure the security of the market protocol is Secure Multiparty Computation [13], which allows a group of parties to execute secure computations without relying on a trusted third party. The communication links between the parties are assumed to be encrypted and synchronous. Security in SMPC is defined by input privacy and protocol correctness. Input privacy requires the input values and calculation results to be kept hidden from other participants unless intentionally revealed, whereas protocol correctness requires that the computations yield the same results as with a trusted third party. The goal is thus to have security no worse than an ideal scheme with a trusted party, i.e., SMPC is allowed to be vulnerable to attacks that also work against the ideal scheme.

The considered adversaries are assumed to be static and honest-but-curious. The static adversary corrupts participants before the start of the protocol but does not corrupt further during execution. Secondly, the honest-but-curious (or passive) adversary seeks only to violate the input privacy property and does otherwise not deviate from the established SMPC protocol [13]. The choice to work with an honest-but-curious adversary is due to the nature of ADMM. While the global stage (iteration) can be implemented with an SMPC secure against misbehaving or active adversaries, these techniques cannot protect against parties misbehaving in the local optimization stage. In particular, an adversary could simply return random values at each iteration and thereby prevent the ADMM from converging without violating any of the SMPC guarantees. Overcoming this issue would necessitate delving into verifiable computing aspects, which are out of the scope.

B. Shamir Secret-sharing Scheme

The SMPC protocol used in this paper is based on the Shamir secret-sharing scheme (SSS) [10]. It provides threshold security based on the number of participants N and a defined threshold $\Theta < N$, where a single, central adversary can corrupt up to Θ participants without compromising the protocol's security. From a high-level perspective, the scheme involves transforming a secret s into a secure *shared* domain, yielding $[s]$, then performing secure calculations on the shared values, and finally, recovering relevant results r, by a reconstruction procedure $[r] \mapsto r$. It is worth noting that due to the inner workings of Shamir secret sharing [10], calculations in an SMPC scheme based on it can only use addition and multiplication operations. As such, the secure market protocol that follows uses a variety of transformations and simplifications to reduce computational complexity as much as possible.

C. Secure Market Protocol Overview

The overview of the proposed market protocol is presented in Fig. 1, with the secure blocks given in green and a dashed line separating the time of market clearing and the time of operation. Each day, market participants initialize the secure market clearing protocol by providing their desired internal parameters such as battery limits, cost coefficients, generation limits, etc. The secure market clearing is then performed, as will be explained in Sec. III-D, which involves the distributed market formulation from Sec. II-D with the updating of global

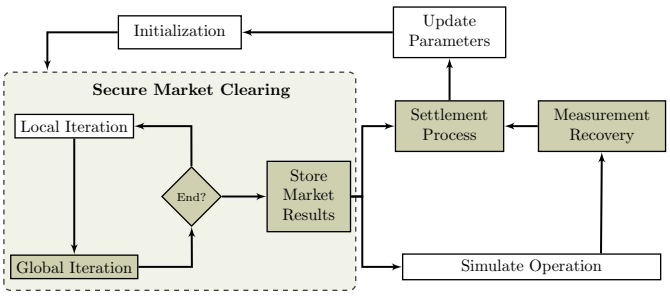

Fig. 1. Overview of the complete (secure) market protocol.

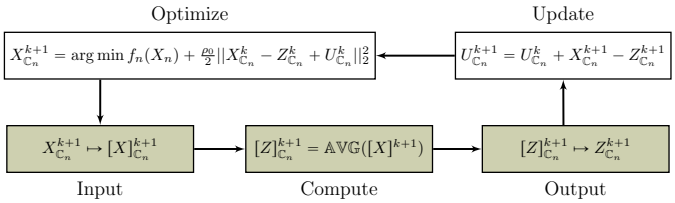

Fig. 2. Illustration of an ADMM iteration with SMPC.

variables and testing for convergence, both implemented securely with SMPC. Upon completion, the parties then store the outcomes of the market for later usage. After the simulated or real market operation, the parties then execute a secure measurement recovery routine (Sec III-E), which ensures the protocol has all the relevant information it needs to proceed. Then, the secure settlement process (Sec. III-F) uses the stored market outcomes and recovered measurements to securely compute the financial balances of all parties. The protocol then ends for the day by having the parties update their internal parameters based on the financial and operational outcomes.

D. Secure Market Clearing

The main usage of SMPC is to secure the distributed optimization process by replacing the central node that performs the coordination calculations (5b). Figure 2 shows how the ADMM update loop from (5) can be secured via SMPC, where white boxes are local operations and boxes shaded green are SMPC operations. Parties hereby securely share the new value of their local variables $X_{\mathbb{C}_n}^{k+1}$, then calculate the shared value of the new global variables $[Z]_{\mathbb{C}}^{k+1}$, and finally, output their true values to the relevant parties for use in local optimization $Z_{\mathbb{C}_n}^{k+1}$. To improve performance, the division operation in the algebraic mean calculation is done locally, meaning the SMPC instance only involves the addition of scaled variables.

The evaluation of convergence criteria is also done within the SMPC instance, which can become computationally expensive due to the non-linear complexity of secure inequality evaluations and Euclidean norm calculations. To address this issue, a stricter version of the residual criteria is used – the infinity norm, instead of the commonly employed ℓ_2-norm.

E. Measurement Recovery

A major concern in digital markets, including the considered LEM, is the "Oracle Problem", where a cyber attack on the underlying measurement systems cause the ground truth and the market's understanding to diverge, causing strain on consumers and network operators [22]. This problem is exacerbated in distribution grids where ensuring complete network observability would require expensive deployment of PMUs or RTUs [23], [24]. The market protocol is thus backstopped with a deterministic and secure measurement recovery procedure to ensure the financial settlement process can be executed even if a subset of nodes fails to report any measurements.

The two main failure modes are nodes not reporting measurements and nodes reporting false measurements [16],

with the latter being difficult to counteract if done at scale. This paper focuses on the former, with the use case being participants purposefully disconnecting their deployed sensor during operation. We assume the use of modified smart meters that can measure the active and reactive nodal injections and the active and reactive incoming line flow at each node. The measurements along the lines are not required since LinDistFlow assumes lossless lines. Under the assumption of altered smart meters and an always reporting substation node, the following recovery procedure generates feasible values for all outstanding net injection measurements, even if they are not guaranteed to report the actual value.

Non-reporting nodes in the network are grouped into disjoint segments, called *islands*, via the UnionFind algorithm [25]. For each such island $\mathcal{Y}$, SMPC is then used to calculate its total active power inflow $\mathrm{P}^{\mathrm{in}}_{\mathcal{Y}}$ and outflow $\mathrm{P}^{\mathrm{out}}_{\mathcal{Y}}$ using the measurements of honest ancestors and descendants. This determines the overall active power net injection of the island, which is then split equally among the members of a given island to ensure burden sharing, as follows

$$\mathrm{net}^{P}_{y} = \frac{\mathrm{P}^{\mathrm{out}}_{\mathcal{Y}} - \mathrm{P}^{\mathrm{in}}_{\mathcal{Y}}}{|\mathcal{Y}|}, \qquad \forall y \in \mathcal{Y}, \tag{8}$$

where $|\mathcal{Y}|$ denotes the cardinality of set $\mathcal{Y}$. Note that the resulting value may not correspond to the withheld value for each $y \in \mathcal{Y}$. Any such deviations are balanced via the balancing mechanism at the time of operation and factor into the settlement process.

F. Secure Settlement Process

The calculation of the financial payoff $\mathcal{B}_n$ for a party n after the market clearing is performed employing the same SMPC protocol used for ADMM coordination. Since the prices used for calculating $\mathcal{B}_n$ are dual variables of constraints, they can also be calculated globally as Lagrange multipliers via SMPC, as explained in [26]. For the final financial balances, the settlement process then factors in the results of the DSO imbalance market (see Section II-E) and the measurement recovery process.

The secure market protocol introduces a break in time between market clearing and final settlement, requiring each node to re-input their own balance $\mathcal{B}_n$ into the SMPC instance computing the final settlement. A Double Verification Scheme (DVS) is therefore proposed to allow secure verification of these input balances by an honest majority of parties, including actors not participating in the protocol, such as the market operator. The DVS combines storing and securely comparing secret-shared balances with the use of cryptographic commitments to verify the self-claimed balances and is inspired by publicly verifiable secret-sharing schemes [27]. Specifically, parties re-share their post-market clearing financial balances, which are securely compared using SMPC to the shared values the other parties stored after market clearing. In the second phase, parties use the perfectly binding ElGamal commitment scheme to "commit" to their balances after market clearing. If needed, they can then "open" the balances in the second phase of DVS to prove that they used the correct value without it actually being revealed [20]. Thus, DVS Phase 1 provides threshold security via SMPC and DVS Phase 2 provides information theoretic security via ElGamal commitments.

DVS thus protects against manipulations by corrupted parties while also providing cryptographic proof for the calculated final financial balances. The secure settlement can then use SMPC to calculate imbalances between the scheduled and actual network values that have both been input securely to find $\mathcal{I}_n$ for each party and then output the final financial balances of all parties.

IV. Results

A. Case Study Setup

The proposed secure market clearing protocol is evaluated on a 16-bus test network depicted in Fig. 3, with the grid-connection being modelled as the network's slack bus. This synthetic case study is constructed using simulated data from various real-world datasets, as described in the following. The network model incorporates nodal demand and solar power generation profiles from an English dataset [28], wind power generation profiles from an Australian dataset [29], and battery sizing based on the Tesla Powerwall [30]. Note that in this specific case study, all prosumers have both a DER and a battery, with their flexible generation stemming entirely from the battery. The power prices are obtained from various sources, including suggested substation tariff [2], wholesale day-ahead prices [31], and regulation market prices [32] from the Danish TSO. Finally, the cost parameters and standard deviations for the prosumers' generation are based on the work in [6], and the battery parameters are derived from [33].

The Python package MPyC [34] was utilized for implementing the SMPC protocol and additionally provided the elliptic curve cryptography necessary for the ElGamal commitment scheme. For the local optimization, the CVXPY optimization library was used with ECOS solver for second-order cone programs and the OCSP solver for quadratic programs. CVXPY also natively supports ADMM. The implementation is done in Python, and the code was executed on a laptop with a 6-core processor and 16 GB of RAM. The implementation of case study, including its SMPC protocols, is publicly available as-is on Github [35].

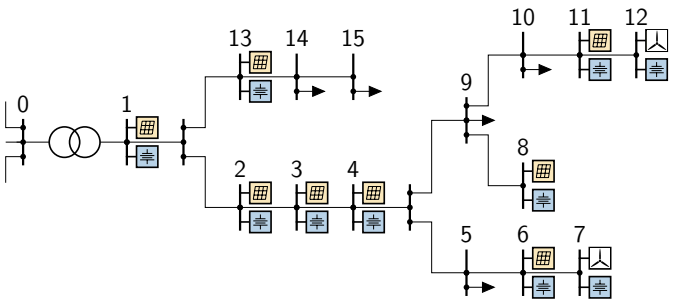

Fig. 3. A synthetic 16-bus test system, with PV generation (in yellow), batteries (in blue), and wind generation (in white).

The results section centers around a comparative analysis of the performance of the following four local market solvers:

- **C-1**: An insecure central solver solving the market problem in (4), with all chance constraints reformulated as second-order cone constraints [6].
- **N-3**: An insecure distributed solver solving the reformulated chance-constrained market problem (4).
- **S-2**: A secure distributed solver solving the reformulated chance-constrained market problem (4) with deterministic voltage magnitude and line flow constraints. Security is imposed using the protocols described in Sec. III.
- **S-3**: A secure distributed solver solving the complete reformulated chance-constrained market problem (4), making it the secure version of N-3.

B. Solver Comparison

The level of security offered by these solvers depends on the extent to which they disclose the private data of the parties during operation. The central solver (C-1), for instance, discloses all parties' information solely to a central computation node. If this central point is compromised, it could potentially threaten input privacy and the correctness of the entire market protocol. On the other hand, in the distributed solvers utilizing ADMM, nodes share certain data only with their immediate neighbors (N-3). Furthermore, when secured with SMPC, widespread data leakage can only occur if a majority of parties are compromised (S-2 and S-3).

Table I compares the four solvers in terms of the time and global iterations required for the algorithm to converge, as well as the resulting relative accuracy. The latter refers to the total active power residual in the LEM relative to the fixed total demand and gauges the feasibility of the solution. All considered solvers converge with similar levels of relative accuracy, but the use of both the decomposition and SMPC extend the time required for the convergence. Notably, N-3 and S-3 require similarly many ADMM iterations, indicating that the calculations in SMPC work adequately. However, comparing S-2 and S-3 shows that full chance constraining doubled the ADMM iterations and tripled the time to convergence. This points to a trade-off between the model complexity and the computation performance. Overall, the results indicate that the secure market protocol via solvers S-2 and S-3 is capable of clearing a complex market with results that are slightly worse but comparable to the centralized solver.

C. Convergence of Secure Market Protocols

As a result of the use of ADMM, all distributed solvers in the paper are guaranteed to eventually converge [26].

TABLE I
COMPARISON OF THE FOUR FEATURED SOLVERS.

Solver	C-1	N-3	S-2	S-3
Time [min]	0.49	66.76	38.06	119.8
Global iterations	1	595	270	593
Rel. accuracy [%]	0	1.8	−1.96	4.21

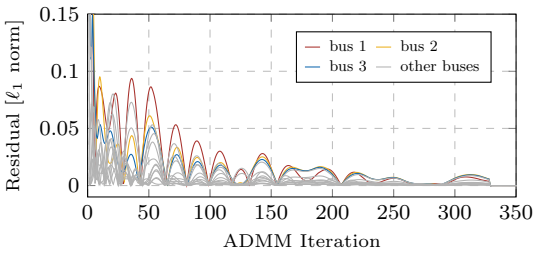

Fig. 4. Voltage magnitude residuals across all nodes when using S-2 solver.

However, as the convergence of the residuals of the voltage magnitudes for the S-2 solver in Fig. 4 highlights, there are some notable dynamics at play. The convergence of ADMM typically exhibits a damped oscillatory behavior [11], as is also evident in the given figure, persisting until approximately iteration 250. Around this iteration, the residuals start to briefly increase again due to the inter-dependency of voltage magnitudes and power flows in the LinDistFlow model. Their constant trade-off (as well as the ripple effect across the nodes) implies that the solver can locally diverge from the expected trends but will eventually recover as the ADMM iterations progress. Interestingly, the residual oscillations are of higher magnitudes at nodes closer to the substation. The observed convergence pattern highlights the possibility for improvement, e.g., through the implementation of adaptive penalty terms or over-relaxation techniques [26], which are, however, beyond the scope of this paper.

D. Market Clearing Outcomes

Figure 5 showcases the different energy balance components over the entire network. It can be observed that the day-ahead energy dispatch of the network is driven by prosumers who use batteries to shift the DER production into the high-demand morning and evening hours. Similar behavior has been seen in previous research [36]. Despite the slight changes in formulations and solution methods, all solvers had very similar energy schedules and prices as the presented S-2 solver, indicating that the secure market protocol is able to clear the market successfully.

The allocation of flexibility participation factors described in (2) by the different solvers is observed to fall into one of the following two arrangements, which is then maintained for the entire day. Specifically, all fully chance-constrained solvers

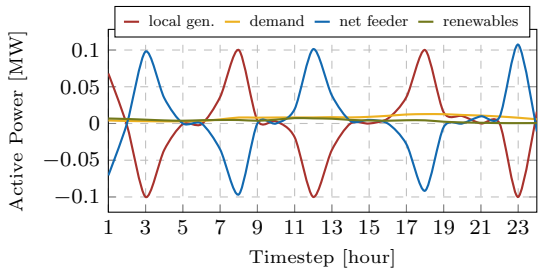

Fig. 5. Global energy balance over a day for the S-2 solver.

TABLE II
FINAL DAY-AHEAD ENERGY BALANCES IN CHF AT SEVERAL NODES FOR SOLVERS N-3, S-2, AND S-3.

Solver	**N-3** (Duals)	**S-2** (Duals)	**S-2** (Sec-Bal)	**S-3** (Duals)	**S-3** (Sec-Bal)
Node 0	−8.76	−7.81	−7.83	−8.41	−8.40
Node 3	0.47	0.43	0.41	0.47	0.47
Node 7	3.71	3.84	3.84	3.71	3.71
Node 15	−0.88	−0.92	−0.92	−0.88	−0.88

assigned the substation the entire flexibility provision, while the S-2 solver, which does not factor in voltage and line flow constraints, splits the flexibility provision equally amongst all nodes with flexible generation. This resulted in the flexibility price being constant throughout the day.

E. Financial Settlement Results

The above-discussed scheduling outcomes are reflected in the post-market clearing financial balances, with Table II comparing the balances of 4 different LEM participants between solvers N-3, S-2, and S-3. Due to the large variability found in flexibility prices across solvers, the table focuses on the day-ahead energy results. As explained in Sec. III-F, in a distributed setting, the market prices can either originate from the dual variables in the local optimization problems that the parties then share with the LEM (Duals) or from global Lagrange multipliers calculated in SMPC (Sec-Bal). To evaluate the impact of using SMPC on these critical calculations, both secure solvers are considered with both price origins and compared to the insecure, distributed solver N-3. The nodes selected for evaluation are the substation node 0, the PV-prosumer node 3, the wind-prosumer node 7, and the pure load at node 15. Across solvers and price origins, there are only minor differences arising from the usage of SMPC, highlighting both SMPC's viability in the LEM and the overall scheme's success at secure market operation.

V. LIMITATIONS & FUTURE WORK

The proof-of-concept nature of the secure and distributed market mechanism proposed in this paper has several limitations regarding its formulation and implementation, which present avenues for future work that are elaborated in the following paragraphs of this section.

The usage of LinDistFlow, a positive-sequence linear network model, follows that of prior works that used it in ADMM-based approaches in distribution networks [24], [12]. It purposefully excludes nonlinear terms and higher-order losses, as described in [26]. This exclusion allows the removal of the branch currents, which would otherwise require bilinear constraints relating them to nodal voltages and branch power flows. This reference further notes that the resulting linearized model introduces only marginal approximation errors. An alternative for future work is the approach of [6], which considers a linear approximation of power flow losses as additional nodal demand. Future work could also include quadratic terms via the Second-Order Cone approach outlined in [6]. Another avenue for future work is extending the approach to multi-phase and unbalanced networks, which would not necessitate major methodological adaptations but rather scaling of the formulation.

Besides avoiding bilinear constraints, the above-mentioned simplification aimed to ensure linear constraints such that distributed optimization would have guaranteed convergence with ADMM [11], [26]. This is desirable due to the limitations of SMPC itself: as its mathematical operations work only with integers, working with decimals necessitates the usage of fixed-point SMPC arithmetic [34]. The resulting SMPC-based ADMM achieves results comparable to the float-based insecure ADMM but nevertheless contains an approximation error. The size of the approximation error can be controlled through the bit size used for fixed-point SMPC, with a trade-off between precision and computational performance of the resulting SMPC protocol. The ADMM-based solution partially counteracts the resulting accumulation of errors, though this relies on the assumption of limited errors. A robust version of the protocol would require the inclusion of error correction techniques, which will be considered in our future work.

Beyond the simplified formulation, the case study used for testing was limited in size and scope. This choice was not made due to the limitations of the formulation or ADMM, as both can be generically scaled to larger network sizes. Instead, it was driven by a performance bottleneck of the simulation environment. The case study involves simulating sixteen separate entities performing local optimization and SMPC operations on a single machine, which reached the memory and processor limits of the simulating compute. The use of either cluster computing or separate physical processors per node would allow for a larger case study, which could be further improved by switching from a high-level language Python library like MPyC [34] to a low-level C++ library like MP-SPDZ [37], which is more suitable for high-performance production usage. To support future work on this matter, we make the source code of our simulation publicly available as-is on Github [35] and refer further to the similarly public code base of prior work this paper draws on [6].

Finally, the case study follows similar approaches in prior works [6], [38] by using a synthetic data set as inputs, created by merging various real-world data sources. These were chosen to simulate a future high-DER distribution network in Europe, with European demand levels and patterns but with wind and solar power levels from warmer climates. The limited case study was in line with the proof-of-concept nature of the paper. Still, the above-mentioned retooling of its simulation would allow future work to simulate more realistic case studies, such as the IEEE European Low Voltage Test Feeder, using a full real-life data set [39].

VI. CONCLUSION

In this paper, we presented a Local Electricity Market framework in a distribution network with uncertain distributed energy resources. To preserve the input privacy of the market participants, we solved the market problem using an

ADMM-based distributed optimization with data exchanges protected by leveraging secure multi-party computation protocols. Amongst the considered secure solvers, S-2 stands out as the recommended choice. It avoids the single point of failure of C-1, has security guarantees via SMPC that N-2 lacks, and achieves comparable results in fewer iterations and minutes than S-3. Note that since the secure market protocol is highly customizable, e.g., in its convergence thresholds, chance constraint bounds, etc., the choice of the preferred solver may vary. Additionally, the protocol may benefit from an increased parallelization or the use of better hardware, which could make S-3 become more suitable.

VII. Acknowledgement

This research was supported by NCCR Automation, a National Centre of Competence in Research, funded by the Swiss National Science Foundation (grant number 51NF40_180545).

References

[1] F. Charbonnier, T. Morstyn, and M. D. McCulloch, "Coordination of resources at the edge of the electricity grid: Systematic review and taxonomy," *Applied Energy*, vol. 318, p. 119188, 2022.

[2] T. Sousa, T. Soares, P. Pinson, F. Moret, T. Baroche, and E. Sorin, "Peer-to-peer and community-based markets: A comprehensive review," *Renewable and Sustainable Energy Rev.*, vol. 104, pp. 367–378, 2019.

[3] J. Kim and Y. Dvorkin, "A P2P-dominant distribution system architecture," *IEEE Trans. on Power Systems*, vol. 35, pp. 2716–2725, 7 2020.

[4] O. Stanojev, Y. Guo, P. Aristidou, and G. Hug, "Multiple ancillary services provision by distributed energy resources in active distribution networks," *arXiv preprint, arXiv:2202.09403v2*, 2022.

[5] A. Ratha, J. Kazempour, A. Virag, and P. Pinson, "Exploring market properties of policy-based reserve procurement for power systems," in *58th Conference on Decision and Control (CDC)*, 2019, pp. 7498–7505.

[6] R. Mieth and Y. Dvorkin, "Distribution electricity pricing under uncertainty," *IEEE Transactions on Power Systems*, vol. 35, no. 3, pp. 2325–2338, 2020.

[7] M. Zugno, T. Jonsson, and P. Pinson, "Trading wind energy on the basis of probabilistic forecasts both of wind generation and of market quantities," *Wind Energy*, vol. 16, pp. 909–926, 2013.

[8] D. K. Molzahn, F. Dörfler, H. Sandberg, S. H. Low, S. Chakrabarti, R. Baldick, and J. Lavaei, "A survey of distributed optimization and control algorithms for electric power systems," *IEEE Transactions on Smart Grid*, vol. 8, no. 6, pp. 2941–2962, 2017.

[9] B. Zhang *et al.*, "An optimal and distributed method for voltage regulation in power distribution systems," *IEEE Transactions on Power Systems*, vol. 30, pp. 1714–1726, 7 2015.

[10] N. Tian, Q. Guo, H. Sun, and X. Zhou, "Fully privacy-preserving distributed optimization in power systems based on secret sharing," *iEnergy*, vol. 1, no. 3, pp. 351–362, 2022.

[11] S. Boyd, N. Parikh, E. Chu, B. Peleato, and J. Eckstein, "Distributed optimization and statistical learning via the alternating direction method of multipliers," *Foundations and Trends in Machine Learning*, vol. 3, pp. 1–122, 2010.

[12] B. Oh, D. H. Lee, W. C. Jeong, and D. Lee, "Distributed optimal power flow for distribution system using second order cone programming and consensus alternating direction method of multipliers," *Journal of Electrical Engineering and Technology*, vol. 17, pp. 999–1008, 3 2022.

[13] U. Maurer, "Secure multi-party computation made simple," *Discrete Applied Mathematics*, vol. 154, pp. 370–381, 2 2006.

[14] T. W. Mak, F. Fioretto, L. Shi, and P. V. Hentenryck, "Privacy-preserving power system obfuscation: A bilevel optimization approach," *IEEE Transactions on Power Systems*, vol. 35, pp. 1627–1637, 2020.

[15] V. Dvorkin, F. Fioretto, P. V. Hentenryck, P. Pinson, and J. Kazempour, "Differentially private optimal power flow for distribution grids," *IEEE Transactions on Power Systems*, vol. 36, no. 3, pp. 2186–2196, 4 2020.

[16] R. Deng *et al.*, "False data injection on state estimation in power systems-attacks, impacts, and defense: A survey," *IEEE Transactions on Industrial Informatics*, vol. 13, pp. 411–423, 4 2017.

[17] D. Pozo, "Linear battery models for power systems analysis," *Electric Power Systems Research*, vol. 212, p. 108565, 4 2022.

[18] C. Gerwin, R. Mieth, and Y. Dvorkin, "Compensation mechanisms for double auctions in peer-to-peer local energy markets," *Current Sustainable/Renewable Energy Reports*, vol. 7, pp. 165–175, 12 2020.

[19] Z. Sui and H. D. Meer, "Bap: A batch and auditable privacy preservation scheme for demand response in smart grids," *IEEE Transactions on Industrial Informatics*, vol. 16, pp. 842–853, 2020.

[20] L. Kamm, "Homological classification of commitment schemes," Master's thesis, University of Tartu, Institute of Computer Science, 2007.

[21] T. Jonsson, P. Pinson, H. A. Nielsen, and H. Madsen, "Exponential smoothing approaches for prediction in real-time electricity markets," *Energies*, vol. 7, pp. 3710–3732, 2014.

[22] G. Caldarelli, "Understanding the blockchain oracle problem: A call for action," *Information (Switzerland)*, vol. 11, pp. 1–19, 11 2020.

[23] K. Dehghanpour, Z. Wang, J. Wang, Y. Yuan, and F. Bu, "A survey on state estimation techniques and challenges in smart distribution systems," *IEEE Transactions on Smart Grid*, vol. 10, pp. 2312–2322, 3 2019.

[24] Y. Wang, Q. Chen, T. Hong, and C. Kang, "Review of smart meter data analytics: Applications, methodologies, and challenges," *IEEE Transactions on Smart Grid*, vol. 10, pp. 3125–3148, 2019.

[25] Z. Galil and G. F. Italiano, "Data structures and algorithms for disjoint set union problems," *ACM Computing Surveys*, vol. 23, no. 3, p. 319–344, 9 1991.

[26] K. Zhang, S. Hanif, and T. Hamacher, "Decentralized voltage support in a competitive energy market," in *2019 IEEE Power & Energy Society Innovative Smart Grid Technologies Conference (ISGT)*, 2019, pp. 1–5.

[27] B. Schoenmakers, "A simple publicly verifiable secret sharing scheme and its application to electronic voting," *Lecture Notes in Computer Science*, vol. 1666, pp. 148–164, 1999.

[28] Northern Powergrid (Northeast) plc, "Dataset (TC5): Enhanced profiling of domestic customers with solar photovoltaics," 8 2014.

[29] J. Dowell, "Australian electricity market operator (AEMO): 5 minute wind power data," 2015.

[30] T. Moore, "How much does the tesla powerwall solar battery cost?" Forbes Home, 3 2023.

[31] Energinet, "Elspot prices," 11 2022.

[32] ——, "Realtime market," 11 2022.

[33] S. Castano, L. Gauchia, E. Voncila, and J. Sanz, "Dynamical modeling procedure of a li-ion battery pack suitable for real-time applications," *Energy Conversion and Management*, vol. 92, pp. 396–405, 3 2015.

[34] B. Schoenmakers, "MPyC–python package for secure multiparty computation," in *Workshop on the Theory and Practice of MPC*, 2018.

[35] M. Franke, "PPDMM." GitHub, 2024. [Online]. Available: https://github.com/Matt-Franke/PSCC_2024_PPDMM

[36] Z. Zhang, R. Li, and F. Li, "A novel peer-to-peer local electricity market for joint trading of energy and uncertainty," *IEEE Transactions on Smart Grid*, vol. 11, pp. 1205–1215, 3 2020.

[37] M. Keller, "MP-SPDZ: A versatile framework for multi-party computation," in *Proceedings of the 2020 ACM SIGSAC conference on computer and communications security*, 2020, pp. 1575–1590. [Online]. Available: https://github.com/mkskeller/mpc-benchmarks

[38] A. Papavasiliou, "Analysis of distribution locational marginal prices," *IEEE Transactions on Smart Grid*, vol. 9, pp. 4872–4882, 9 2018.

[39] T. Morstyn and M. D. McCulloch, "Multiclass energy management for peer-to-peer energy trading driven by prosumer preferences," *IEEE Transactions on Power Systems*, vol. 34, pp. 4005–4014, 9 2019.