

Blockchain for EUPHEMIA Market Transparency[†]

Tsvetomir Gospodinov , Mariana Atanasova  and Eliza Stefanova 

Faculty of Mathematics and Informatics, Sofia University “St. Kliment Ohridski”, 5 James Bourchier Blvd., 1164 Sofia, Bulgaria; matanasova@fmi.uni-sofia.bg (M.A.); eliza@fmi.uni-sofia.bg (E.S.)

* Correspondence: cbgospodin@fmi.uni-sofia.bg

[†] Presented at the 15th International Scientific Conference TechSys 2026—Engineering, Technologies and Systems, Plovdiv, Bulgaria, 14–16 May 2026.

Abstract

EUPHEMIA, the Pan-European day-ahead electricity market-coupling algorithm, operates in a centralized manner that restricts independent auditability and has been characterized as pseudo-transparent. We propose a blockchain-based architecture that improves the transparency and verifiability of the market-coupling process while preserving participant confidentiality. It combines off-chain computation with selective on-chain publication and treats the three principal data categories of the EUPHEMIA pipeline separately: order books, network constraints, and clearing outputs. Zero-knowledge proofs utilizing zk-STARKs are employed to verify the integrity of the order book aggregation process and specific network-constraint sub-processes, whereas Merkle commitments ensure tamper-evident anchoring of publicly disclosed data. zk-STARKs are selected over CRS-based alternatives to eliminate the trusted-setup governance overhead associated with EUPHEMIA’s multi-jurisdictional structure. The estimated AIR trace size reaches approximately 2,225,000 rows in the worst-case NEMO (EPEX SPOT) scenario. A correction proof generated by the Regional Coordination Centre (RCC) requires an AIR of 641 trace rows when a binding network constraint exceeds its threshold during the review of Transmission System Operator (TSO) submissions. This trace size corresponds to an estimated proving time of 1 to 30 s based on reported STARK prover throughput. End-to-end verification of welfare maximization remains infeasible due to the lack of a complete public algorithm specification. Preliminary calibrated estimates are provided, and full empirical benchmarking remains as future work.

Keywords: EUPHEMIA; blockchain; zero-knowledge proofs; zk-STARKs; single day-ahead coupling



Academic Editors: Nikola G. Shakev,
Valyo N. Nikolov, Nikolay R.
Kakanakov and Sevil A.
Ahmed-Shieva

Published: 16 July 2026

Copyright: © 2026 by the authors.
Licensee MDPI, Basel, Switzerland.
This article is an open access article
distributed under the terms and
conditions of the [Creative Commons
Attribution \(CC BY\)](https://creativecommons.org/licenses/by/4.0/) license.

1. Introduction

EUPHEMIA [1], the Pan-European Hybrid Electricity Market Integration Algorithm, has computed prices and energy allocations for European day-ahead and intraday markets since February 2014. It solves a welfare-maximizing optimization subject to cross-border transfer capacities and network constraints set by Transmission System Operators (TSOs). Inputs are the buy and sell orders submitted to Nominated Electricity Market Operators (NEMOs) by its participants. Each TSO computes cross-zonal capacity constraints from its grid model and submits them to a Regional Coordination Centre (RCC), which validates them against neighboring constraints and returns corrections if required. The TSO then forwards the final constraints to the Market Coupling Operator (MCO), which executes EUPHEMIA. Market clearing then yields electricity prices, order acceptance decisions, net positions per bidding zone, and inter-zonal flows. The process offers limited transparency.

Participants and external auditors cannot independently verify the inputs, intermediate computations, or final outcomes.

We propose a blockchain-based verification framework for the EUPHEMIA market-coupling process with the following contributions:

- An analysis identifies the confidentiality and integrity challenges in the EUPHEMIA market-clearing process, including attack vectors and trust assumptions.
- A blockchain-based architecture supports transparent and auditable EUPHEMIA execution through on-chain commitments and zero-knowledge proofs that enable independent verification and preserve confidentiality.
- A hybrid model combines off-chain computation with on-chain commitments and verification, accommodating the algorithm's computational demands within the throughput constraints of current blockchain platforms.
- A zk-STARK proof framework addresses selected correctness properties of order aggregation, network-constraint corrections, and clearing outputs without disclosing sensitive data, avoiding trusted setup ceremonies and relying on standard cryptographic assumptions.
- A feasibility analysis estimates AIR trace size, prover work, verifier cost, on-chain storage requirements, and compatibility with day-ahead market timing constraints.

2. Background

2.1. EUPHEMIA Market Coupling and CACM Framework

Three categories of data pass through the EUPHEMIA pipeline (Figure 1): order books, network constraints, and clearing outputs. Market participants submit individual orders to NEMOs, which aggregate them into supply and demand curves and forward these, together with block, complex, and merit orders, to EUPHEMIA. Each TSO computes cross-zonal capacity constraints from its grid model and submits them to its RCC for regional coordination. The RCC proposes corrections informed by neighboring grids, and the TSO then submits the finalized constraints to the MCO. Only aggregated parameters, including Available Transfer Capacity (ATC) and Net Transfer Capacity (NTC) values, are publicly disclosed, while the underlying constraints and grid models remain confidential. The MCO executes EUPHEMIA and produces the clearing output, which is returned to NEMOs and published as the market result. Each category carries distinct confidentiality requirements and integrity verification challenges, and each is anchored on-chain through a different strategy.

A distributed ledger combines immutability, public auditability, and decentralized governance within a single infrastructure layer. Append-only logs or notary schemes provide similar guarantees under narrower trust assumptions, but a distributed ledger fits the multi-jurisdictional CACM framework, in which no single institution is universally trusted.

2.2. Zero-Knowledge Proof Systems

Verifiable computation has its roots in the early work on interactive proof systems and zero-knowledge proofs by Goldwasser, Micali, and Rackoff [2]. Over time, these theoretical foundations led to practical systems. For example, Pinocchio [3] made it possible to efficiently verify outsourced computations—regardless of their size—in about 10 milliseconds. As research progressed, newer schemes sought different balances between proof size, verification speed, and setup requirements. Groth16 [4] produced constant-size proofs with quick verification, but required a separate trusted setup for each circuit. Later, PLONK [5] introduced a universal setup that could be reused for many circuits. More recently, zk-STARKs [6] eliminated the need for any trusted setup by using only collision-resistant hash functions, though this comes with the trade-off of larger proofs.




Figure 1. EUPHEMIA data flow diagram.

3. Related Work

Blockchain-based designs for auction and double-auction markets combine cryptographic commitments with on-chain verification of settlement outcomes [7]. Energy-sector work has examined privacy-preserving market clearing, mainly in local energy markets and peer-to-peer trading platforms [8,9]. These systems typically coordinate distributed energy resources or support decentralized trading while preserving the confidentiality of participant bids.

To our knowledge, this is the first work that combines large-scale centralized market coupling, confidentiality constraints, and public verifiability via zk-STARKs. In that setting, transparency must coexist with bid confidentiality, and published clearing outcomes must be independently verifiable.

The Austrian Energy Agency report [10] critiques the current transparency framework and introduces the concept of “pseudo-transparency”: information is formally published yet remains effectively inaccessible due to the disproportionate effort required to locate or process it. It identifies three shortcomings: the opacity of the aggregation process, the lack of verifiable input data, and the absence of independent audit mechanisms.

Most prior work on transparency in electricity markets focuses on regulatory oversight and market efficiency [11], treating auditability and data accessibility as conditions for trust in market outcomes while recognizing the need to protect commercially sensitive information. Current European market-coupling mechanisms publish data *ex post* but provide no cryptographic guarantees of correctness.

Recent advances in verifiable computation, particularly zero-knowledge proof systems, allow verification of complex computations without disclosing private inputs [6,12]. Researchers have applied zk-SNARK and zk-STARK systems to verifiable machine learning and outsourced computation, but applications to large-scale optimization problems with mixed constraints remain limited.

Existing approaches either focus on decentralized market designs with limited scalability or on transparency mechanisms without formal correctness guarantees. None address the combination considered here: large-scale centralized optimization, confidentiality, and public verification of market outcomes.

4. System and Adversary Model

The architecture assumes participants may not fully trust the entities that execute the market-clearing process. Required guarantees cover manipulation of order data, incorrect aggregation of bids, and misreporting of market outcomes. Four classes of adversary motivate this scope:

- Nominated Electricity Market Operators (NEMOs) collect bids from market participants and submit aggregated order books to the market-coupling algorithm. A malicious or compromised NEMO could exclude valid orders, alter bid parameters, or publish incorrect aggregate curves.
- Transmission System Operators (TSOs) supply transmission capacity data and network-related constraints that directly influence the feasible market outcome. A malicious or compromised TSO could misreport cross-zonal capacities, omit constraints, or submit inconsistent technical data to influence the clearing result.
- Market Coupling Operator (MCO) executes EUPHEMIA and may publish results that do not correspond to the actual inputs or that deviate from the prescribed clearing logic.
- Independent observers, including competitors and other market participants, may attempt to infer confidential bidding strategies from publicly available data or observable market outputs.

Attack vectors against the integrity of clearing include unauthorized modification or omission of submitted orders before aggregation, publication of aggregated curves that do not accurately reflect the underlying bids, misreported transmission capacities or network constraints, and market outcomes inconsistent with the actual execution of the clearing algorithm.

All adversaries are computationally bounded and cannot break the collision resistance of the employed hash functions or the soundness of the zero-knowledge proof system.

5. Proposed Architecture

Selected EUPHEMIA inputs and outputs are published on-chain, while zero-knowledge proofs verify aggregation without disclosing confidential data.

Each NEMO records a commitment to its submitted orders before clearing begins. Each TSO records a commitment to its network constraints. The MCO then records a commitment to the outcome and a zero-knowledge proof for the checked correctness properties once clearing completes.

A public blockchain record gives market participants and regulators a tamper-evident audit path for the clearing process. The computational demands of the algorithm and the throughput constraints of current blockchain platforms prevent direct on-chain execution. Computation remains off-chain, with selective on-chain publication.

The choice between a public and a permissioned ledger involves distinct trade-offs: public chains have stronger censorship resistance at higher operational cost, whereas permissioned deployments reduce cost but reintroduce a consortium-level trust assumption.

Three verification flows make up the design, as shown in Figure 2. The market-orders flow requires each NEMO to publish a Merkle root for its order book. An AIR attests that the aggregated curves were derived from the committed orders. The network-constraints flow combines TSO inputs and RCC review. Each TSO commits the resulting aggregated parameters on-chain, and the RCC publishes a correction ZKP whenever a corrective capacity adjustment is triggered by a binding constraint exceeding a threshold. The market-clearing flow runs off-chain. The MCO executes EUPHEMIA on these inputs, then commits the clearing results together with a ZKP for flow feasibility, order consistency, and non-negative volumes.




Figure 2. System architecture overview.

6. Proof Construction

6.1. Selection of Proof System

Most ZKP systems require a Common Reference String (CRS), a set of public parameters upon which both proof generation and verification depend. The CRS is typically produced through a trusted setup in the form of a multi-party computation (MPC) ceremony, in which each participant contributes randomness to the generation of the public parameters and subsequently destroys their individual secret contribution. The security guarantee rests on the assumption that at least one participant behaves honestly. The resulting CRS is sound under that condition and cannot be exploited to forge proofs.

The EUPHEMIA coupling algorithm is updated at least annually, and the ceremony must be repeated each time. Coordination is complex: 16 NEMOs, 40 TSOs, a rotating Market Coupling Operator, national regulators, and thousands of market participants are involved, with no single entity universally trusted by all others. The choice of ceremony participants is both a technical and political challenge. The process must also ensure that every party genuinely destroys its secret material afterward.

zk-STARKs [6] derive their security entirely from collision-resistant hash functions and public randomness. They eliminate the need for a common reference string and, with it, the trusted parameter-generation ceremony and its associated governance burden. On-chain verification has polylogarithmic verifier complexity, although zk-STARKs currently incur larger proof sizes and higher prover costs than CRS-based schemes including Groth16 [4] or PLONK [5]. The removal of the trusted setup justifies this trade-off, given EUPHEMIA's multi-jurisdictional governance structure and the annual recurrence of any ceremony that would otherwise be required. The design adopts zk-STARKs as the preferred proof scheme.

6.2. Market Orders

Supply and demand curve aggregation is a fundamental transparency challenge. Verification requires access to underlying individual orders, which are not publicly available. Consequently, neither market participants nor regulatory authorities can independently verify integrity.

Zero-knowledge proofs (ZKPs) resolve the trade-off between verifiability and confidentiality. A prover can demonstrate correct aggregation without disclosing underlying data, which supports regulatory compliance and preserves the ability to audit computations.

The verification process relies on a preliminary commitment step. Each NEMO publishes a Merkle root committing to its complete set of orders. The Merkle tree uses a STARK-friendly hash, known as Poseidon [13], which reduces AIR trace rows per hash

compared to alternatives including SHA-256. The published root acts as a tamper-evident snapshot of the order book. Any modification to an order invalidates the root and becomes detectable at verification.

The AIR expresses the aggregation pipeline in four stages (Figure 3). Commitment checks ensure that private orders match the published Merkle root. Range checks enforce correct assignment of orders to price steps. Aggregation sums quantities per price step across all participating NEMOs. Equality checks ensure consistency with the publicly declared aggregated curve.

The four operations form a sequential pipeline in which each stage depends on its predecessor. A verified proof attests to the integrity of the aggregation process, from order submission to published output.

Complex, block, and merit orders are transmitted without aggregation. Their Merkle roots can be published on-chain without zero-knowledge proofs.




Figure 3. Market orders flow.

6.3. Network Constraints

Network constraints define physical and regulatory limits on cross-border electricity flows. Each TSO computes constraints from its grid model, security limits, and remedial actions, and submits them to its RCC for regional coordination. The RCC proposes corrections informed by neighboring grids, and the TSO then submits the finalized constraints to the MCO for inclusion in the clearing optimization.

The constraints are not fully public. Only aggregated parameters such as ATC and NTC values are disclosed, while the underlying detailed constraints and grid models remain confidential. The aggregated values can be anchored on-chain via Merkle commitments, with ZKPs attesting to correct derivation from the underlying constraints without exposing them. The MCO retrieves these committed inputs as tamper-evident inputs to EUPHEMIA.

The detailed IGMs underlying these constraints, by contrast, are not disclosed to market participants. ZKP-based verification of IGMs is not currently feasible due to their

complexity. IGMs include busbar topology, power flows, contingencies, and remedial actions. The corresponding AIR representations have execution-trace sizes that exceed what is computationally feasible under current ZKP technology. Strict CACM timing constraints further prevent computationally intensive proof generation.

Zero-knowledge proofs nevertheless remain applicable to targeted sub-processes. A zk-STARK generated by the RCC can attest that a corrective capacity adjustment was triggered by a binding constraint exceeding a threshold, without disclosing the identity or parameters of the congested network element.

The AIR comprises four operations, as illustrated in Figure 4. Data binding fixes the witness and ties it to upstream commitments through two Poseidon hashes. The combined hash over capacity and threshold locks the prover to a fixed witness, preventing post hoc substitution. The capacity-only hash binds the witness to the value previously committed on-chain by the originating TSO, preventing arbitrary capacity values. A range check verifies that the capacity lies within the valid operational domain. A comparison confirms that the capacity exceeds the predefined security threshold. An inequality assertion links the private witness to the public claim.




Figure 4. Network constraints flow.

The on-chain record of each correction includes the proof, the witness commitment, the correction claim, and a pointer to the originating TSO commitment. The threshold and the operational range bounds are resolved through the governance configuration of the deployment, either by reference to a governance contract or by embedding the values directly in the transaction payload when the configuration is held off-chain. These accompanying inputs allow any verifier to reconstruct the public arguments required for proof verification without consulting private grid data.

6.4. Market Clearing

The computational complexity of EUPHEMIA prevents direct on-chain execution. The algorithm is executed off-chain, with its inputs and outputs committed to the blockchain to ensure an immutable record.

Market clearing results are publicly disclosed, but their correctness is not self-evident. A consistency check against network constraints and order data would require re-execution of the algorithm, which remains computationally infeasible for external observers.

Figure 5 summarizes the clearing flow. Independent verification requires the Market Coupling Operator (MCO) to generate a zk-STARK proof attesting to core correctness properties of the clearing outcome. These properties include the feasibility of cross-zonal flows with respect to published network constraints, the consistency of accepted orders with the committed order books, and the non-negativity of cleared volumes. On-chain verification of the proof gives any party public auditability and does not require re-execution of the algorithm or disclosure of confidential inputs.




Figure 5. Market clearing flow.

An alternative approach relies on executing the clearing algorithm within a Trusted Execution Environment (TEE), which has hardware-enforced isolation and remote attestation. A cryptographically signed attestation confirms that the prescribed algorithm was executed on verified enclave hardware. The TEE approach reduces computational overhead but introduces additional trust assumptions in hardware vendors and has weaker public verifiability guarantees than zero-knowledge proofs.

7. Security Analysis

A central security question is whether published proofs can be checked without exposing private market data. Four standard properties of verifiable computation and cryptographic protocols [2,6] guide the analysis: soundness, zero-knowledge, tamper-evidence, and public verifiability. These properties address auditability and aggregation-opacity concerns raised about the current market-coupling framework [10].

7.1. Soundness

Soundness means a computationally bounded adversary cannot convince verifiers of a false statement. The most direct example here is a NEMO attempting to publish an aggregated curve that does not match the committed order book. A verifying proof gives cryptographic assurance that the checked aggregation statement is consistent with the committed data, subject to the soundness of the underlying FRI-based STARK proof system [6].

The concrete soundness error is not fixed by the architecture alone. It depends on implementation parameters, including the code rate, blowup factor, query count, number of FRI rounds, and field size. A larger number of FRI queries reduces the chance that an invalid execution trace passes the verifier's low-degree and consistency checks. A production deployment would need to publish its concrete proof parameters together with the intended security target, for example 128-bit soundness. Soundness here relies on the standard guarantees of the selected STARK construction. We do not claim a fixed security level for a particular parameter set.

7.2. Zero-Knowledge

Zero-knowledge addresses the main confidentiality tension in the design. Market participants need assurance that aggregation was performed correctly. Individual orders and their attributes must remain private. The proof should reveal no information beyond the public statement being checked: the relevant Merkle roots and aggregated curves.

zk-STARKs support this property by masking the execution trace with fresh randomness before committing to the low-degree extension [6]. The values opened during verification do not directly expose the private witness. The Fiat–Shamir transformation and Merkle commitments are instantiated through hash functions in practice. The intended hash family is Poseidon because it is designed for use inside zero-knowledge proof systems [13].

Poseidon must still be instantiated carefully. Its parameters must match the target field, state width, capacity, and security level. The Poseidon design provides parameter-selection rules and cryptanalysis for algebraic and statistical attacks in zero-knowledge proof systems [13].

7.3. Tamper Evidence

Cryptographic commitments combined with an append-only ledger provide tamper evidence. A NEMO's commitment binds its order book, so any later modification fails to match the published commitment and becomes visible.

Each NEMO publishes a Poseidon-based Merkle root for its order book on the blockchain ledger, as described in Section 6.2. Collision resistance of the hash function ensures that a computationally bounded adversary cannot feasibly find a different order book with the same root. Any change, removal, or replacement of an order after commitment invalidates the root, and any party checking the ledger can detect it.

The ledger adds a second protection layer. A commitment that has reached the confirmation depth required by the chosen chain can be rewritten only by violating that ledger's consensus and settlement assumptions. Two mechanisms provide this protection: cryptographic binding through the Merkle root and operational immutability through ledger finality. The combination allows regulators, market participants, and independent auditors to check that proofs refer to the data originally committed.

7.4. Public Verifiability

Public verifiability allows independent audit. Any party with access to the ledger can retrieve the published proofs and verify them without participating in the original computation and without seeing confidential order data. Verification in the order-aggregation flow requires only the public inputs, including Merkle roots and aggregated curves, together with the zk-STARK proof.

The verifier complexity is $\mathcal{O}(C \cdot \log^2(\beta N))$ field operations, where C is the number of transition constraints and N is the trace length [6]. The order-aggregation flow considered in this paper has $C \sim 42$ constraints and a padded LDE domain of 2^{23} , yielding an esti-

mated verifier cost of approximately 22,000 field operations. The verifier cost is essentially independent of the number of orders.

7.5. Guarantees and Assumptions

A valid proof for the order-aggregation flow attests to the checked aggregation relation between the committed order book and the published aggregate curve. Individual orders remain private during verification. Tamper evidence binds the NEMO to the order book snapshot committed before clearing.

The guarantees hold under standard cryptographic assumptions: adversaries are computationally bounded, Poseidon is instantiated with parameters appropriate for the chosen field and target security level, and the non-interactive STARK proof is analyzed in the random oracle model after applying the Fiat–Shamir transformation.

There is also an engineering assumption. The underlying blockchain must reach the required level of consensus-level finality before commitments are treated as immutable.

8. Feasibility Evaluation

The feasibility analysis uses four criteria: correctness, transparency, confidentiality, and on-chain feasibility. Correctness concerns the relation checked by the proof. Transparency concerns external verification of the published data. Confidentiality concerns protection of private inputs. On-chain feasibility concerns whether verification and storage fit realistic blockchain constraints.

8.1. AIR Complexity: Market Orders

Proof-generation overhead is proportional to the number of orders processed by each NEMO. The reported daily average of 4828 orders [14] is never processed by a single instance because proof generation is distributed. Each NEMO proves only its own order book. The largest NEMO determines the worst-case AIR size. Per-NEMO order volumes are not publicly reported. Traded volume per bidding zone is a reasonable proxy because order count correlates with market activity.

The CACM Annual Report [14] identifies EPEX SPOT as the largest NEMO by traded volume, with 654 TWh [15] (47.2% of annual volume), followed by Nord Pool with 518 TWh [16] (31.3%). The two NEMOs together account for 78.5% of total annual traded volume. Taken as proxies for order concentration, these shares give a worst-case per-NEMO order count of approximately $4828 \times 0.472 \approx 2279$ orders for EPEX SPOT.

Table 1 presents the worst-case AIR trace size estimate. zk-STARK prover complexity scales as $\mathcal{O}(N \log^2 N)$ in the trace size N [6], so the trace of approximately 2,225,000 rows yields $2,225,000 \times 441 \approx 9.8 \times 10^8$ operations, where $\log^2(2,225,000) \approx 441$.

Table 1. Estimated AIR trace size for worst-case NEMO (EPEX SPOT, ~2279 orders).

Operation	Count	Rows Each	Total Rows
Order commitment	8191 hashes	~250	2,047,750
Price-step assignment	2279	~70	159,530
Summation across step	2279	~1	2279
Equality checks	17,000	~1	17,000
Total			~2,225,000

Current STARK prover implementations report throughput ranging from approximately 10^5 to 10^7 trace rows per second depending on field choice, hardware configuration, and degree of parallelism. At blowup factor $\beta = 4$ in the Goldilocks field with width $w \approx 10$, the worst-case AIR requires an estimated 1 to 30 s of proving time on a modern multi-core

server. The estimate combines the AIR trace size in Table 1 with reported throughput ranges from existing STARK prover implementations and public prover frameworks, including Stone [17].

Proof generation is distributed. Each NEMO proves only its own order book, and NEMOs can generate these proofs in parallel. The day-ahead market process leaves approximately 42 min between gate closure and result publication. The estimate is compatible with the CACM timing constraint. Proof generation is therefore not the system bottleneck under realistic CACM timelines.

Table 2 summarizes the estimated prover working memory. Prover memory is dominated by the LDE and composition polynomial. Estimates use the Goldilocks field (8 bytes per element) with a degree-2 extension for FRI queries, blowup factor $\beta = 4$, 64 FRI queries ($64 \times \log_2 \beta = 128$ bits of provable soundness), maximum remainder polynomial degree 31, and no grinding. The padded trace length is 2^{22} rows, the LDE domain is 2^{24} rows, and the effective AIR width is $w \approx 10$ columns.

Table 2. Estimated prover working memory at 128-bit provable security.

Component	Size Formula	Estimate
Padded trace	$2^{22} \times 10 \times 8 \text{ B}$	0.34 GB
LDE	$2^{24} \times 10 \times 8 \text{ B}$	1.34 GB
Composition LDE	$1.5\text{--}2 \times \text{LDE}$	2.0–2.7 GB
Merkle tree	$2^{24} \times 32 \text{ B} + \text{internal nodes}$	1.0 GB
FRI in-place	$\sim 0.5 \times \text{LDE}$	0.7 GB
Auxiliary	—	0.5 GB
Total		$\sim 6\text{--}7 \text{ GB}$

The estimate fits a single 64-core x86_64 server with 32 GB of RAM. GPU acceleration is optional, and the 42-min day-ahead clearing budget covers the estimate with margin.

8.2. AIR Complexity: Network Constraints

Table 3 summarizes the corresponding trace size estimate. The trace of $N \approx 641$ rows yields $\mathcal{O}(N \log^2 N) \approx 641 \times \log^2(641) \approx 56,000$ operations of prover work, which is negligible compared to the Order Book AIR. Verification requires $\mathcal{O}(\log^2 N) \approx 87$ operations. The witness contains the capacity value, the threshold, and a binding to the upstream TSO commitment. Verification does not require access to the full grid model.

The RCC correction proof reuses the order book parameter set: $\beta = 4$, Goldilocks field, $w \approx 10$, and 128-bit provable security. The padded trace length is 2^{10} rows and the LDE domain is 2^{12} rows. Working memory is dominated by the LDE ($2^{12} \times 10 \times 8 \text{ B} \approx 0.3 \text{ MB}$) and composition polynomial, totaling on the order of a few megabytes. Proving time stays below 100 milliseconds. Any commodity server runs the RCC correction proof comfortably, well inside the CACM review window.

Table 3. Estimated AIR trace size for RCC correction proof.

Operation	Count	Rows Each	Total Rows
Data binding	2 hashes	~ 250	500
Range check	1	~ 70	70
Threshold comparison	1	~ 70	70
Inequality assertion	1	~ 1	1
Total			~ 641

8.3. Storage Requirements: Market Clearing

The EUPHEMIA clearing process yields day-ahead prices, net positions, and cross-zonal flows. A configuration with 17 bidding zones and 96 market time units (MTUs) requires 1632 values each for prices and net positions, plus approximately 3360 values for cross-zonal flows. At 64 bits per value, the data require roughly 53 KB per clearing session. The total output size, including acceptance results for Block, Complex, and Merit Orders, is approximately 200 KB.

Direct publication is possible but expensive. A more practical alternative commits to the full output through a Merkle tree, which lets verifiers check individual elements through inclusion proofs while reducing on-chain storage.

The zk-STARK proof generated by the MCO verifies a subset of correctness properties: cross-zonal flows satisfy the published network constraints, accepted orders are consistent with the committed order books, and cleared volumes are non-negative.

The framework trades additional proof and commitment costs for public verifiability of selected correctness properties.

The proof does not establish welfare maximization. The absence of a complete public EUPHEMIA specification prevents formal encoding of the objective function and all associated constraints into a verifiable AIR. Verification remains partial: it addresses feasibility and consistency rather than optimality.

9. Discussion and Limitations

The prover work estimates rely on the theoretical $\mathcal{O}(N \log^2 N)$ bound. Actual proving times depend on the implementation.

The verification scope covers feasibility, non-negativity, and acceptance consistency. Welfare maximization remains outside the scope of this paper because the full EUPHEMIA specification is not public. Verification of the Individual Grid Model remains infeasible at present. Three implementation choices also stay open: ledger type, multi-party approval process, and AIR update procedure. The institutions operating the market resolve them.

The current EUPHEMIA process has near-zero cryptographic overhead but provides no independent verifiability. The proposed approach introduces a 2–3% latency increase and additional compute costs in exchange for public auditability. Order book aggregation, network-constraint handling, and clearing outputs no longer depend on operator trust alone. Each step becomes verifiable through cryptography. On-chain Merkle commitments provide a tamper-evident audit trail that ex-post publication cannot match.

The costs are modest. Worst-case NEMO proof generation takes 1 to 30 s and needs 6 to 7 GB of memory. RCC correction proofs complete in under 100 milliseconds. Each round adds a few hundred kilobytes of on-chain storage. Added end-to-end latency stays at 2 to 3% of the 42-min day-ahead window, and verifier work sits below typical L1 blockchain throughput. The overhead is small relative to the gain in cryptographic verifiability across the three data flows.

A practical deployment path could begin with optional post-trade verification, allowing operators to adopt the system incrementally without modifying the core clearing process.

10. Future Work

An implemented STARK prover for the order book aggregation circuit is the immediate next step, with measured proving time, verifier cost, and peak memory use on representative hardware.

End-to-end welfare-maximization verification remains the largest open problem. A practical path may combine partial zk-STARK proofs over disclosed sub-procedures with

auditor-side replay of the published EUPHEMIA executable, once a full algorithm specification exists.

The approach extends to the XBID intraday market, where shorter trading cycles leave less time for proof generation and demand smaller AIRs. Recursive proof composition is worth examining, because it may allow modular verification of Individual Grid Model (IGM) sub-AIRs without generating a single proof for the full network model.

Reducing trace size by employing alternative arithmetization-friendly hash primitives, lookup arguments, and batched commitments could decrease prover costs. Governance and key-management frameworks for the NEMO, TSO, and MCO roles require further development.

11. Conclusions

We presented a blockchain-based architecture that makes EUPHEMIA market coupling auditable while preserving order and network confidentiality. The design publishes commitments, selected outputs, and zero-knowledge proofs on-chain, keeping sensitive inputs off-chain.

zk-STARKs fit this setting because they avoid trusted setup and support public verification under EUPHEMIA's multi-jurisdictional governance. Estimated AIR sizes are 2,225,000 rows for worst-case NEMO order aggregation (EPEX SPOT) and 641 rows for an RCC correction proof. The corresponding proving time of 1 to 30 s fits within the 27-min residual window after EUPHEMIA execution. These figures are calibrated estimates, not empirical measurements.

The proposed framework is intended to assist market regulators, NEMOs, RCCs, and external auditors requiring independent verification of clearing processes. At present, the framework addresses Market Order Aggregation and selected RCC sub-processes. End-to-end welfare-maximization verification remains unattainable without public disclosure of the complete algorithm. Prover implementation and end-to-end timing measurements remain as future work.

Author Contributions: Conceptualization, T.G.; methodology, T.G. and M.A.; formal analysis, T.G.; investigation, T.G.; writing—original draft preparation, T.G.; writing—review and editing, M.A. and E.S.; supervision, M.A. and E.S. All authors have read and agreed to the published version of the manuscript.

Funding: This research received no external funding.

Institutional Review Board Statement: Not applicable.

Informed Consent Statement: Not applicable.

Data Availability Statement: Data are contained within the article.

Conflicts of Interest: The authors declare no conflicts of interest.

Abbreviations

The following abbreviations are used in this manuscript:

AIR	Algebraic Intermediate Representation
ATC	Available Transfer Capacity
CACM	Capacity Allocation and Congestion Management
CRS	Common Reference String
EUPHEMIA	Pan-European Hybrid Electricity Market Integration Algorithm
FRI	Fast Reed-Solomon Interactive Oracle Proof of Proximity
IGM	Individual Grid Model

LDE	Low-Degree Extension
MCO	Market Coupling Operator
MPC	Multi-Party Computation
NEMO	Nominated Electricity Market Operator
NTC	Net Transfer Capacity
RCC	Regional Coordination Centre
TEE	Trusted Execution Environment
TSO	Transmission System Operator
XBID	Cross-Border Intraday
ZKP	Zero-Knowledge Proof
zk-SNARK	Zero-Knowledge Succinct Non-Interactive Argument of Knowledge
zk-STARK	Zero-Knowledge Scalable Transparent Argument of Knowledge

References

1. All NEMO Committee. *EUPHEMIA Public Description: Single Price Coupling Algorithm*; All NEMO Committee: Brussels, Belgium, 2025. Available online: <https://www.nemo-committee.eu/assets/files/euphemia-public-description.pdf> (accessed on 5 February 2026).
2. Goldwasser, S.; Micali, S.; Rackoff, C. The knowledge complexity of interactive proof systems. *SIAM J. Comput.* **1989**, *18*, 186–208. [CrossRef]
3. Parno, B.; Howell, J.; Gentry, C.; Raykova, M. Pinocchio: Nearly practical verifiable computation. In *Proceedings of the IEEE Symposium on Security and Privacy (S&P)*; IEEE: Piscataway, NJ, USA, 2013; pp. 238–252. [CrossRef]
4. Groth, J. On the Size of Pairing-Based Non-Interactive Arguments. Cryptology ePrint Archive. Paper 2016/260. 2016. Available online: <https://eprint.iacr.org/2016/260> (accessed on 20 February 2026).
5. Gabizon, A.; Williamson, Z.J.; Ciobotaru, O.-M. PLONK: Permutations over Lagrange-Bases for Oecumenical Noninteractive Arguments of Knowledge. Cryptology ePrint Archive. Paper 2019/953. 2019. Available online: <https://eprint.iacr.org/2019/953> (accessed on 20 February 2026).
6. Ben-Sasson, E.; Bentov, I.; Horesh, Y.; Riabzev, M. Scalable, Transparent, and Post-Quantum Secure Computational Integrity. Cryptology ePrint Archive. Paper 2018/046. 2018. Available online: <https://eprint.iacr.org/2018/046> (accessed on 20 February 2026).
7. Jia, X.; Wang, L.; Cheng, K.; Jing, P.; Song, X. A blockchain-based privacy-preserving and collusion-resistant scheme (PPCR) for double auctions. *Digit. Commun. Netw.* **2023**, *11*, 116–125. [CrossRef]
8. Erdayandi, K.; Mustafa, M.A. PP-LEM: Efficient and privacy-preserving clearance mechanism for local energy markets. *Sustain. Energy Grids Netw.* **2024**, *39*, 101477. [CrossRef]
9. Mengelkamp, E.; Gärtner, J.; Rock, K.; Kessler, S.; Orsini, L.; Weinhardt, C. Designing microgrid energy markets: A case study: The Brooklyn Microgrid. *Appl. Energy* **2018**, *210*, 870–880. [CrossRef]
10. Holzmann, A.; Zwieb, L.; Knaus, K.; Harrucksteiner, A.; Sahin, A. *Increasing Transparency for Flow-Based Market Coupling in European Electricity Trading: Barriers, Solutions, and Key Indicators*; Austrian Energy Agency: Vienna, Austria, 2020. Available online: https://oesterreichsenergie.at/fileadmin/user_upload/IncreasingTransparencyForFBMC_20201216.pdf (accessed on 15 February 2026).
11. European Commission. Commission Regulation (EU) 2015/1222 of 24 July 2015 establishing a guideline on capacity allocation and congestion management. *Off. J. Eur. Union* **2015**, *L197*, 24–72. Available online: <https://eur-lex.europa.eu/eli/reg/2015/1222> (accessed on 1 March 2026).
12. Chiesa, A.; Hu, Y.; Maller, M.; Mishra, P.; Vesely, P.; Ward, N. Marlin: Preprocessing zkSNARKs with Universal and Updatable SRS. Cryptology ePrint Archive. Paper 2019/1047. 2019. Available online: <https://eprint.iacr.org/2019/1047> (accessed on 20 February 2026).
13. Grassi, L.; Khovratovich, D.; Rechberger, C.; Roy, A.; Schofnegger, M. Poseidon: A new hash function for zero-knowledge proof systems. In *Proceedings of the USENIX Security Symposium*, Online, 11–13 August 2021; pp. 519–535.
14. All NEMO Committee. *CACM Annual Report on the Single Day-Ahead and Intraday Market Coupling in Europe 2024*; All NEMO Committee: Brussels, Belgium, 2025. Available online: <https://www.nemo-committee.eu/assets/files/cacm-annual-report-2024.pdf> (accessed on 25 February 2026).
15. EPEX SPOT. *Annual Trading Results of 2024: Power Trading on EPEX SPOT Reaches All-Time High*; EPEX SPOT: Paris, France, 2025. Available online: <https://www.epexspot.com/en/news/annual-trading-results-2024-power-trading-epex-spot-reaches-all-time-high> (accessed on 18 February 2026).

16. Nord Pool. *Nord Pool Reports Encouraging Growth in 2024*; Nord Pool: Oslo, Norway, 2025. Available online: <https://www.nordpoolgroup.com/en/message-center-container/newsroom/exchange-message-list/2025/q1/nord-pool-reports-encouraging-growth-in-2024/> (accessed on 18 February 2026).
17. StarkWare. Stone: STARK One: STARK Prover and Verifier. GitHub Repository. 2023. Available online: <https://github.com/starkware-libs/stone-prover> (accessed on 20 February 2026).

Disclaimer/Publisher’s Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.