electronics

MDPI

Article

A Decentralized Framework to Gather and Certify Green Energy Data in Demand Response Programs

Daniele Marletta, Alessandro Midolo and Emiliano Tramontana *

Dipartimento di Matematica e Informatica, University of Catania, 95125 Catania, Italy; daniele.marletta@phd.unict.it (D.M.); alessandro.midolo@unict.it (A.M.)
* Correspondence: tramontana@dmi.unict.it; Tel.: +39-095-7383008

Abstract

The increasing adoption of renewable energy sources introduces significant variability in power generation, requiring effective strategies to ensure maintain grid stability. Incentive-based demand response programs provide a practical solution for balancing supply and demand, however disputes may arise over energy data integrity. The existing solutions frequently rely on centralized authorities, exposing a single point of failure, or high costs and privacy limitation of recording granular data on-chain. To address this challenge, we propose a decentralized framework that separates cloud storage from integrity certification. This system employs a community aggregator to collect high-frequency energy measurements, store the raw data in the cloud, while anchors unique cryptographic hashes for batch of raw data to a public blockchain. This process creates an auditable and tamper-evident record of data. By recording only hashes on chain, our approach achieves privacy and scalability. Evaluation using a real-world Australian dataset confirms that the system enables transparent dispute resolution, with blockchain transaction costs consistently representing less than 0.10% of the total incentives awarded to participants.

Keywords: energy systems; distributed energy resources; renewable energy sources; prosumers; demand response; data certification; data integrity; smart contracts

1. Introduction

The adoption of renewable energy sources (RESs) has experienced steady growth, driven by the need to reduce dependence on fossil fuels. Solar photovoltaic (PV) power generation, in particular, has emerged as one of the key sustainable technologies, with an average annual growth rate exceeding 35% over the past decade [1,2]. Global PV capacity almost doubled between 2019 and 2022, and the overall share of RESs in the electricity mix is projected to increase from 30% in 2023 to 46% by 2030 [3,4]. Despite this rapid growth, RESs are characterized by intrinsic variability, as their generation depends on meteorological conditions. This unpredictability creates significant challenges for energy systems, including the optimal placement and detection of solar panels [5,6], as well as ensuring grid stability and balancing net energy [7].

In recent years, Demand Response (DR) programs have experienced a growing interest as a methodology to tackle the unpredictable power generation output given by RESs [8]. Broadly, there are two main categories of DR strategies: (i) price-based and (ii) incentive-based [9]. Price-based approaches employ dynamic energy pricing to encourage end users to modify their consumption patterns. However, studies have shown that the efficacy of this approach can be limited [10]. Conversely, incentive-based DR programs offer

Academic Editor: Herminio Martínez-García

Received: 3 April 2026
Revised: 16 June 2026
Accepted: 17 June 2026
Published: 19 June 2026

rebates to consumers who agree to reduce their load demand in response to signals from system operators. However, when financial incentives are awarded, disputes may arise. For example, a customer's claim can contest the reward amount and may request higher compensation. Establishing proof of data integrity is essential to deterring tampering activities and resolving remuneration conflicts. In this context, we propose an approach that leverages the distributed blockchain network to foster trust among all the entities participating in an energy system of this kind. Specifically, data hashes are stored on-chain to: (i) provide proof of integrity for the associated data, and (ii) detect data tampering that results from a hash mismatch.

Blockchain technology has been widely used to certify data and replace centralized time stamping systems [11–14]. In particular, many approaches in the literature have proposed the integration of public blockchain with DR programs to enhance and accelerate their adoption [15–17]. However, these approaches exhibit limitations such as high costs, significant delays in finalizing large numbers of transactions, and a lack of user data privacy. A critical challenge in applying blockchain to energy systems lies in ensuring scalability and energy efficiency. Public blockchains, while providing strong guarantees of immutability and transparency, often suffer from high transaction costs, limited throughput, and non-negligible energy consumption due to consensus mechanisms. These limitations are particularly relevant in DR scenarios, where large volumes of fine-grained energy data are generated at high frequency. To address these issues, our approach is explicitly designed to minimize on-chain operations through data batching and off-chain storage, thereby reducing both computational overhead and energy consumption associated with blockchain interactions. By limiting blockchain usage to compact cryptographic representations of data, the proposed framework achieves a scalable and energy-efficient certification process suitable for real-world, large-scale deployments.

Our decentralized software system provides proof of integrity for energy measurements sent by smart meters during DR events. In our architecture, the key actor is the community aggregator, which continuously collects energy data, assembles them into batches of predetermined size, and requests the writing of batch hashes on the public blockchain. Once finalized, the receipt of the writing operation is associated with the original data, which are subsequently committed to a remote cloud server for long-term storage. This dual-layered storage strategy ensures verifiability, scalability, and security while minimizing the interaction costs with the blockchain [18–20]. The integrity of the recorded energy data can be verified by comparing the hash of the data stored in the cloud with the corresponding hash retrieved from the blockchain [21]. To evaluate the economic feasibility of our approach, we conducted extensive experiments using a publicly available dataset of energy data released by the Australian government [22]. The results of our analysis reveal that our approach provides data certification while achieving optimal blockchain interaction and cost efficiency.

The key contributions of this work are: the design and implementation of a decentralized system that employs the immutability property of blockchain storage to provide proof of integrity for recorded energy data while also ensuring data privacy; a batching strategy that significantly reduces blockchain costs while preserving data integrity; a validation for the approach using real-world data and provide a detailed cost analysis.

The remainder of this paper is structured as follows. Section 2 surveys the most relevant approaches in the literature and compares them with our proposed system. Section 3 presents the proposed software architecture for certifying DR events. Section 4 describes our experiments conducted on a real-world DR dataset. Section 5 discusses the results. Finally, Section 6 draws our conclusions.

2. Related Works

Research on ensuring data integrity and authenticity has evolved from cryptographic timestamping schemes to blockchain-based certification frameworks, and more recently to their integration within energy and demand–response (DR) systems. Traditional Long-Term Time-Stamping (LTTS) schemes rely on cryptographic primitives, trusted intermediaries, or hierarchical infrastructures to guarantee verifiable records over time [23]. Meng et al. [24] proposed a formal security model for LTTS using Message Authentication Codes (MACs), archival services, and transient key management, offering strong theoretical guarantees. However, these schemes remain anchored to trusted authorities, creating single points of failure (and high storage and key-maintenance overheads) that our decentralized approach eliminates. Similarly, Vigil et al. [25] explored decentralized timestamping for digital archives, yet still relied on timestamping servers, leaving the system partially centralized and vulnerable to single points of failure. In contrast, our approach eliminates trusted intermediaries entirely by leveraging the immutability of a public blockchain to provide verifiable timestamps and data integrity proofs without centralized control, while maintaining low operational complexity.

Blockchain has emerged as a natural foundation for replacing centralized timestamping systems. Meng and Chen [26] proposed to store hash digests on a private blockchain to ensure verifiability without third-party intervention, while Wilson [12] developed a permitted blockchain solution integrating off-chain storage. Although these frameworks improved autonomy, they introduced governance centralization and weakened transparency guarantees. The utility of blockchain for ensuring transactional finality is further underscored in other high-stakes sectors; for instance, Agarwal et al. [27] demonstrated its application in financial market clearing and settlement systems to enhance transparency and reduce the need for manual reconciliation. Similarly, Zhang et al. [28] engaged Ethereum for timestamping, achieving immutability but at the cost of privacy exposure and high transaction fees. More recent solutions, such as IoETTS by Qian et al. [29] and the accumulator-based design of Zhang et al. [14], attempted to improve scalability and cost-efficiency but either suffered from high energy consumption due to consensus mechanisms or lacked empirical validation. Compared to these, our dual-storage architecture attains the same verifiability at significantly lower cost and latency. By storing only batched data hashes on the public blockchain and committing full data sets to cloud storage, our method minimizes blockchain transactions while preserving transparency and integrity—making it well-suited for high-frequency data in energy systems.

Parallel to these works, blockchain has been increasingly explored to increase transparency, security, and coordination in energy management and DR programs. Koukaras et al. [15] offered a comprehensive review of blockchain integration in smart grids, identifying potential factors for improved DR efficiency and trust, but acknowledging open issues related to scalability, interoperability, and privacy. To address coordination challenges, Bâra and Oprea [30] proposed a Digital Twin model to synchronize energy community activities, though such models primarily focus on operational visibility rather than immutable data certification. Lucas et al. [31] demonstrated a Hyperledger Fabric-based DR registry for secure validation of service delivery, achieving low latency and resource usage, yet their design remains limited to small-scale, permitted networks and lacks mechanisms for large-scale aggregation and certification of measurement data. Moreover, numerous DR optimization frameworks [32–36] have addressed the operational and economic coordination of DR participants, including the use of multi-agent reinforcement learning for reliability assessment in electricity market simulations [37]. These models generally assume trustworthy data exchange and focus on cost minimization, ignoring the

need for tamper-proof certification or automated dispute resolution mechanisms between aggregators, consumers, and operators.

By unifying blockchain-based immutability with off-chain storage efficiency, our framework addresses the specific certification needs of high-frequency energy data. Unlike existing blockchain–DR solutions that focus primarily on energy trading or price mechanisms, we target the integrity and verifiability of smart meter data used for incentive-based DR participation. The proposed framework introduces a community aggregator that collects energy measurements, batches them, stores their hashes on a public blockchain, and archives the raw data on a remote cloud server. This hybrid strategy provides verifiable proof of integrity, scalability, and privacy preservation, while drastically reducing blockchain costs compared to prior schemes that commit raw or frequent data on-chain. Furthermore, by coupling blockchain immutability with off-chain efficiency, our approach enables transparent and cost-effective dispute resolution in incentive-based DR programs, which is a feature absent from previous LTTS, blockchain timestamping, or DR optimization studies.

A comprehensive comparison between our proposed framework and the discussed literature is presented in Table 1. As illustrated, while existing LTTS schemes ensure data authenticity, they typically compromise on decentralization. Conversely, standard blockchain timestamping and financial settlement approaches [27,28] provide immutability however often overlook privacy and scalability—two critical factors for smart meter data. Furthermore, although recent energy-specific models using Digital Twins or Multi-Agent Reinforcement Learning [30,37] offer advanced coordination and market simulation, they lack the tamper-proof certification required for secure, dispute-free incentive-based DR. By integrating off-chain storage with batched blockchain-based integrity proofs, our approach unifies these research directions into a practical, economically feasible, and privacy-preserving system that ensures verifiable energy data integrity for real-world DR operations.

Table 1. Comparison of Related Works vs. Proposed Framework.The ✗ symbol indicates that the approach does not satisfy the corresponding property, while the ✓ symbol indicates that it does. The ∼ symbol indicates that the approach partially satisfies the corresponding property, or that its satisfaction depends on specific conditions.

Approach	Decentral.	Integrity	Scalable	Energy/DR	Privacy	Off-Chain
Traditional LTTS [23,24]	✗	✓	✗	✗	✓	✓
BC Timestamping [12,28]	✓	✓	✗	✗	✗	∼
Financial Clearing [27]	✓	✓	✗	✗	∼	✗
BC for DR Registry [31]	∼	✓	∼	✓	✗	✗
Digital Twin/Sim. [30,37]	✗	✗	✓	✓	∼	✓
DR Optimization [32,33]	✗	✗	✓	✓	✗	✗
Proposed Framework	✓	✓	✓	✓	✓	✓

3. System Architecture

At a high level, the system consists of three main components: (i) smart meters that collect data, (ii) an aggregator that processes and batches data, and (iii) a blockchain layer that certifies data integrity.

3.1. Energy Communities

Energy Communities (ECs) were introduced by European Union regulations as voluntary organizations of citizens involved in energy-related activities [38]. As an evolution of the traditional consumer, the new prosumer was introduced to describe end users who can also generate electricity using RESs, mainly solar PV and energy storage systems (ESS) [39]. PV prosumers can inject surplus energy into the electrical grid, while ESS-enabled prosumers can offer flexibility services by strategically charging or discharging their batteries [40]. The aggregator is an entity that remotely manages the data of EC

members. It acts as a central orchestrator of distributed energy resources and coordinates energy optimization strategies to support grid operations [41]. Figure 1 shows the main entities operating within energy communities.

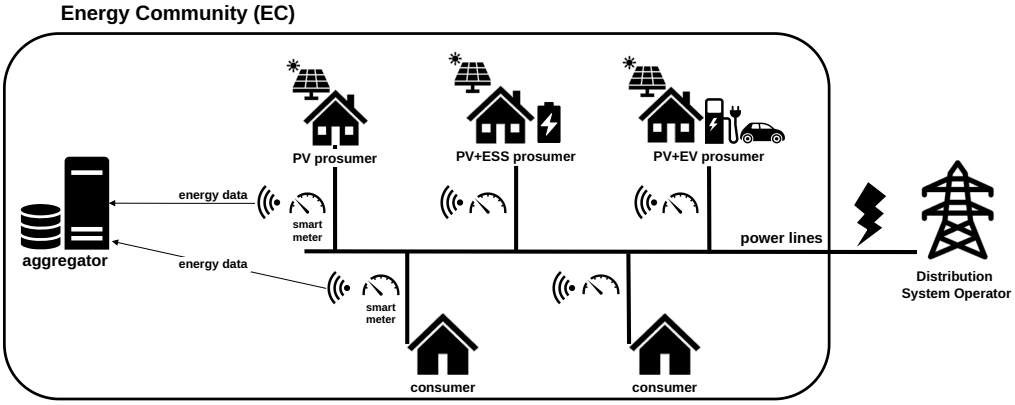

Figure 1. Overview of the residential Energy Community (EC). Prosumers are equipped with several types of distributed energy resources, including solar PV panels, battery energy storage systems (BESS), and electric vehicles (EVs). Energy data are periodically transmitted to an aggregator by a dedicated smart metering network.

The distribution system operator (DSO) is responsible for managing the power distribution network and infrastructure, and delivers electricity to end users. In our system, we consider two types of energy users: (i) traditional consumers and (ii) prosumers, who can both consume and produce electricity. Prosumers own one or more distributed energy resources from RESs, including rooftop solar PV panels, battery energy storage systems (BESS) and electric vehicles (EV). Each user in the EC is equipped with a smart meter, an embedded IoT device that measures the incoming or outgoing flow of electrical power. SMs periodically send energy data over the Internet to the aggregator, which stores them in a local database. The proposed architecture assumes the existence of secure communication channels between the smart meters and the aggregator. While device-level authentication is a critical component of initial data collection, our research focuses on the subsequent risk of data modification, ensuring the integrity of the data once it has reached the aggregator and is moved to permanent storage.

The aggregator is an agent that collects and handles energy data for all community members. Using data transmitted by SMs, the aggregator continuously monitors the EC's net energy balance and forecasts short-term energy profiles for the community. Anomalies and peaks are identified by comparing observed and predicted values and predetermined thresholds, which are defined to ensure compliance with optimal energy profiles or to promote grid stability in agreement with the DSO. When these thresholds are exceeded, the aggregator requests users to modify their energy usage pattern to mitigate or altogether prevent undesirable deviations from the baseline profile. For example, if monitoring data indicate a situation of overconsumption, users are notified to reduce demand. Additionally, ESS-enabled and vehicle-to-grid (V2G) prosumers are notified to discharge their batteries into the electricity network. Conversely, when net energy exceeds the predetermined production threshold, prosumers are prompted to use the energy surplus to charge their batteries. To improve member compliance with energy usage modification requests, the aggregator can issue an incentive-based Demand Response (DR) program, where users are rewarded for reducing their energy load. However, the distribution of remuneration inherently introduces the risk of disputes, such as a user contesting their reward amount or claiming eligibility for a higher reward. To address this challenge, system operators must

implement mechanisms to ensure that incentives are correctly assigned and protect against data manipulation.

In an energy community, the participants can trust the aggregator because it operates within a highly transparent, tightly regulated, and collaborative ecosystem. European Union Directives and national regulatory frameworks formalize the duties of the aggregator, ensuring they act as consumer-aligned demand service providers bound by strict legal compliance and exchange contracts. This institutional oversight guarantees fair value distribution, data privacy, and transparent management of local energy clusters. While a cloud-based aggregator is not entirely immune to risks, its deployment on a cloud server makes failure due to cyberattacks or hardware unavailability highly unlikely. Modern cloud infrastructure relies on a distributed architecture that utilizes robust redundancy, automated failover mechanisms, ensuring that if one server goes down, another seamlessly takes its place. Furthermore, automated DDoS mitigation, advanced encryption standards, and real-time defence mechanisms ensure a highly resilient environment for the energy community's system.

3.2. Demand Response Program and Certification of Energy Data

Blockchain provides a decentralized and tamper-resistant mechanism for data certification. In essence, blockchain acts as a distributed ledger that leverages cryptographic functions to securely record data in a chain of blocks. Blockchain properties ensure that the content in a finalized block remains immutable, preventing deletion and modification. Therefore, blockchain fosters trust for entities that do not rely on a central authority and require integrity for the stored data. Ethereum and Ethereum Virtual Machine (EVM)-compatible blockchains have further developed this technology by enabling the definition of programmable Smart Contracts [42]. The execution of these contracts is triggered by a set of predetermined conditions, and their source code is committed to and stored on the blockchain, ensuring transparency. We leveraged the properties of blockchain technology to design an innovative framework that ensures secure and transparent certification of energy data during DR programs. The proposed system certifies critical energy data, promoting trust and accountability. Figure 2 shows the architecture of the proposed system, highlighting the interaction between the key entities and the main data flow of our proposed system.

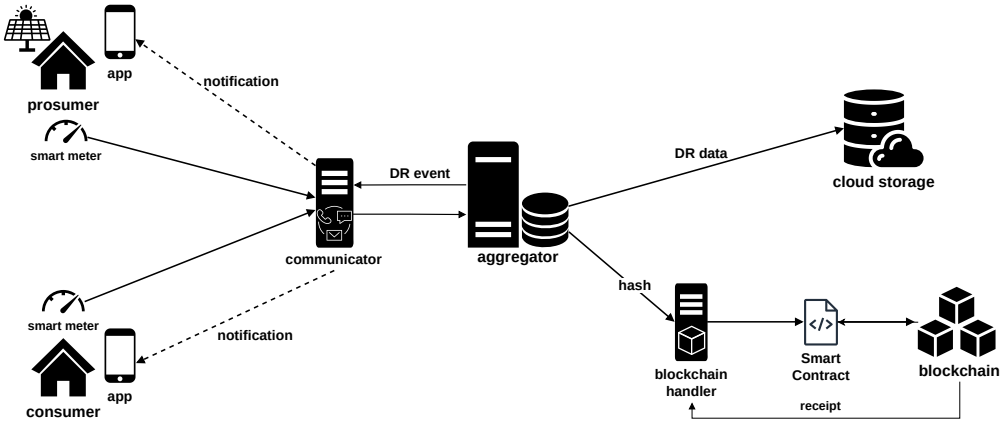

Figure 2. Architecture diagram illustrating the integration of blockchain technology in Energy Communities (EC) for certifying energy data in Demand Response (DR) programs. The aggregator utilizes a communicator to interact with EC members, requests the recording of data hashes on the blockchain via a dedicated handler, and sends DR data to a cloud component for persistent storage.

The main agent in our proposed system is the EC aggregator, which is responsible for starting DR events, managing energy data sent from smart meters, and orchestrating

a dual-layer storage strategy that certifies data integrity. The aggregator continuously monitors the community's energy usage and employs predefined thresholds to ensure that the net energy profile does not diverge from optimal patterns. When the aggregator detects or forecasts that the reference thresholds are exceeded, it defines and activates a new DR event. DR parameters set by the aggregator include: (i) the start time, (ii) the event duration, and (iii) the financial incentive per unit of reduction awarded to complying users. This information is transmitted to the communicator service, which is responsible for delivering DR event data to EC members' smartphones using messages, SMS or in-app notifications, depending on the users' preferences. Each user independently evaluates the conditions and rewards associated with the event and responds to the communicator with the decision to participate or not in the program.

The communicator sets a timeout immediately after notifying users, and, since participation to DR programs is on a voluntary basis, it is assumed that users who do not reply within the predefined time interval will not take part in the DR event. Consequently, unless users explicitly opt-in for the event, either by sending a reply or configuring the app, their energy data are not evaluated, and they will not receive a reward even if they met the reduction target during the event. However, as long as the DR event is still active (i.e., before its scheduled end time), a user can decide to opt-in by notifying the communicator. In this situation, user eligibility for the financial reward is evaluated only for energy data transmitted after their reply and until the end of the event.

After collecting user replies, the communicator forwards them to the aggregator, which determines the subset of EC members who agreed to modify their energy usage patterns in exchange for a financial reward. Consequently, the aggregator updates the DR event record in its local database to set up the energy data certification procedure for all participating users. Smart meters send timestamped energy measurements to the communicator, which validates and forwards them to the aggregator. The aggregator applies a filtering operation to the collected energy data to extract measurements of participating users, which are inserted into the DR event record. The process of collecting energy measurements and storing them in the local database continues until the end of the DR event. Therefore, the aggregator acts as a reliable short-term storage system, ensuring that all DR event records are securely kept in its local database until the event ends.

The aggregator uses blockchain technology to provide a transparent and immutable record of hashes for DR-related energy data. This blockchain integration enables the certification of data integrity and therefore renders data tampering ineffective. The aggregator performs the integrity certification process for every user who has opted into the event, including those who fail to achieve a reduction in demand. Recording data for non-complying participants is essential for dispute resolution, as it provides immutable proof that their energy usage did not meet the criteria for a financial reward. This comprehensive audit trail ensures that all claims, whether regarding successful participation or eligibility denials, can be verified against the blockchain record.

When the number of data points reaches a predetermined threshold, the aggregator assembles them in a data batch and forwards them to the blockchain module. This module is responsible for computing the corresponding hash and ensuring that the hash is correctly stored on the blockchain for immutable storage by invoking a Smart Contract. Once finalized, the receipt of the blockchain writing operation is associated with the original data, which are eventually committed to a remote cloud server for long-term storage. Figure 3 illustrates the sequence diagram for the entire DR event workflow.

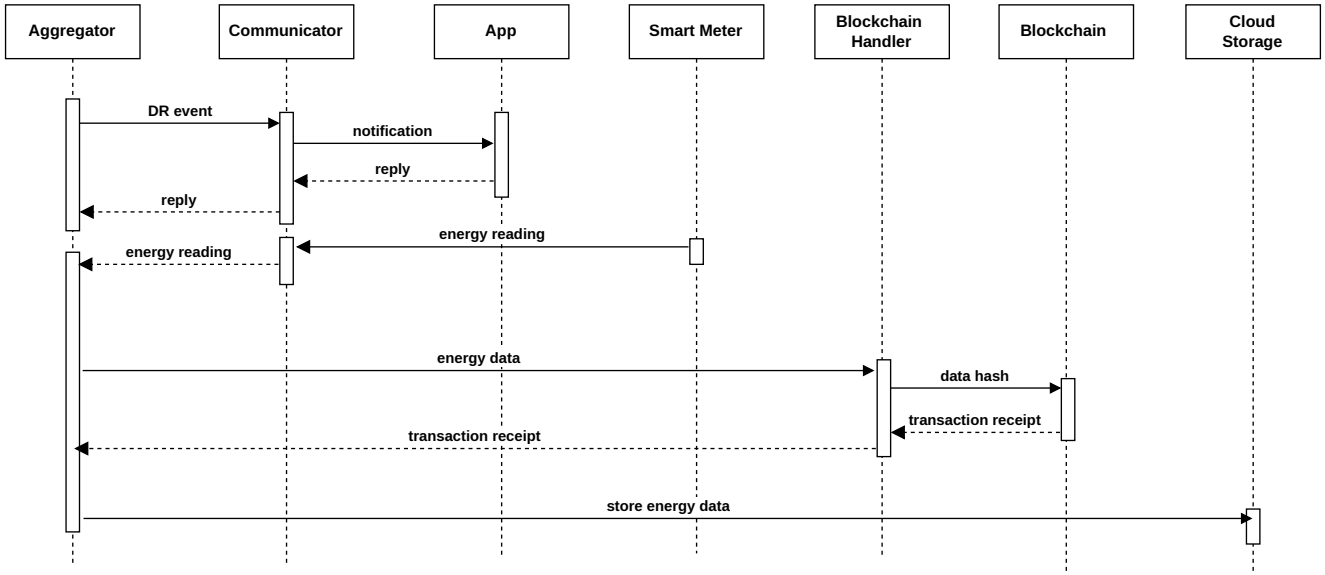

Figure 3. Sequence diagram illustrating data interactions within the proposed architecture. The communicator notifies users of DR events, while the aggregator collects energy data and instructs the blockchain handler to record the data hash on the distributed ledger. After obtaining the receipt, DR data are sent to cloud storage.

The hashes serve as proof of integrity for the corresponding data or, if a mismatch occurs between the hash of data in the storage and the corresponding hash retrieved from the blockchain, they reveal that the storage data were modified or are incorrect. Consequently, malicious data tampering activities on the stored data become futile, as they will be revealed by comparing the hashes. If a technical fault prevents the finalization of the transaction, the aggregator requests the blockchain handler module to retry the hash writing operation. In these situations, the energy data remain in the local database and are not sent to the remote cloud storage until the receipt is generated. This approach ensures that each data batch is always associated with an immutable record, even in the presence of temporary faults. This strategy of integrating traditional cloud storage with the properties of blockchain achieves an optimal balance between system scalability and security. The immutability property ensured by blockchain technology provides a reliable proof of data integrity, promoting trust between all entities in the energy system. Moreover, by storing only the hash in the distributed public ledger, the privacy of user data is ensured. Any attempt to forge data on the blockchain would require compromising the aggregator's credentials or the consensus of the public ledger itself, which represents a significantly higher barrier than tampering with a local database. By anchoring hashes to a public blockchain, we ensure that the state of the data as received by the aggregator remains immutable. This methodology establishes a robust security framework for Demand Response programs where the fundamental risk involves the retrospective alteration of energy data to manipulate financial settlements. The proposed approach ensures that once data are recorded, any attempt to modify historical values for financial gain is detectable through the blockchain audit trail.

To provide strong security, the SHA-256 algorithm is used to compute hashes, which represent a unique digital fingerprint for the corresponding data. By storing the hashes in the blockchain, a public, transparent and immutable record is established, which can be used by stakeholders and authorities to verify that the original data have not been altered. In our proposed approach, each traceable data entry consists of the following information: the ID of the DR event, the user ID, the energy measure, and the associated timestamp.

Since storing a hash for each individual data point could result in high transaction costs for using the blockchain, data are aggregated in batches based on predetermined thresholds. Subsequently, the hash function is applied to the data batch and the output is sent to an EVM-compatible blockchain for immutable storage. Hashes can later be retrieved from the distributed ledger to confirm the integrity of the corresponding energy data maintained in the remote cloud storage.

3.3. Smart Contract DataHashStorage

The DataHashStorage smart contract defines and exposes the two key functions for writing and reading hashes on the blockchain: storeDataHash(bytes32 _dataHash) and getBlockTimestampFromHash(bytes32 _dataHash). The blockchain module executes the storeDataHash function, which receives the hash of a data batch as input. This transaction results in saving the hash in a mapping defined by the contract, which associates each hash with the timestamp of the block containing the transaction. The block timestamp represents the time interval in seconds between the Unix epoch and the minting of the block. After the transaction is completed, the corresponding receipt is returned to the aggregator, concluding the data batch processing. Periodically, the aggregator sends all data for which there is a confirmed hash registration to the remote cloud server for permanent storage.

The getBlockTimestampFromHash() function is used to obtain proof of integrity for a specific data batch. The data hash is provided as input for the function, which accesses the smart contract's internal mapping to recover the corresponding block timestamp. The existence of such a timestamp proves that the hash is recorded on the blockchain, thus validating the original data batch. Conversely, if no match is found in the internal mapping for the provided hash, the function returns the default value of zero. This proves that the provided input hash was not stored on the blockchain and therefore the corresponding data batch was altered.

To ensure scalability in broader deployments involving multiple Energy Communities, each aggregator can be associated with a unique smart contract and a dedicated blockchain address. This architectural separation prevents data conflicts and hash collisions across different communities. On the blockchain layer, concurrent writing requests from multiple aggregators are managed by the network's inherent sequential transaction processing and the aggregator's internal buffering system, which prevents bottlenecks during high-activity periods.

3.4. Resolving Disputes with a Cost-Effective Solution

Disputes regarding energy data and corresponding financial remuneration may occur when an EC member claims different values for their energy usage during the DR event compared to the ones that are stored in the system. When such situations arise, data hashes retrieved from the blockchain enable us to verify or disprove the user's claim. This verification process unfolds as follows: firstly, the complete energy data batch is recovered from the cloud storage and the corresponding hash is determined. Subsequently, a second hash is computed for data that include the disputed energy values and timestamps. Finally, the getBlockTimestampFromHash() function is invoked to confirm the presence of these hashes on the blockchain storage. Successful hash retrieval proves the integrity and validity of the corresponding data, while failure to confirm the existence of the hash demonstrates its incorrectness, thus concluding the dispute.

Notably, our proposed approach uses a public blockchain, enabling independent data validation by authorized entities. When facing legal or regulatory proceedings, it is critical to ensure to external stakeholders that recorded data have remained unaltered and are not compromised by unauthorized modifications. The immutability property of the blockchain

ensures that successfully stored hashes cannot be tampered with, providing a secure and verifiable audit trail. The usage of a public blockchain, as opposed to an Ethereum sidechain or a private blockchain or network controlled by the aggregator, provides a trustless and decentralized record resistant to centralized tampering.

This dual-layer storage strategy represents a scalable and cost-effective solution, providing several advantages. Storing on blockchain only hashes derived from one-way cryptographic functions ensures that private information is not exposed on the public distributed ledger, thereby enforcing compliance with privacy regulations. Moreover, the complete set of data is securely maintained in the cloud storage (off-chain), which can only be accessed by authorized entities. Furthermore, the data validation process is streamlined, requiring only the retrieval of the corresponding immutable hash from the blockchain and its comparison with the hash derived from the data from the cloud. An exact hash match proves the integrity of the original data, confirming that they were not altered by malicious entities who obtained access to the cloud storage. Finally, a blockchain transaction involving only a 32-byte hash leads to a significant cost decrease compared to storing a large amount of energy data records on-chain.

Smart meters typically transmit energy measurements at regular time intervals (e.g., every 30 min). To achieve an optimal balance between data granularity and blockchain costs, the aggregator uses a set of predetermined thresholds to determine the frequency of hash writing transactions. When the total amount of collected measures reaches a threshold, the aggregator assembles them into a data batch and requests the blockchain module to store its hash and return the receipt. The thresholds are defined to maintain a consistent size for the data batch, regardless of the number of participants in the DR event, which can differ from event to event. The next section presents the experimental assessment of our proposed approach, evaluated using a real-world energy dataset.

To ensure robust security against malicious actors, the framework implements an explicit tamper-detection pipeline. In the event that raw energy metrics stored within the cloud tier are altered, appended, or prematurely deleted, the off-chain cryptographic hash computed during an audit will fundamentally diverge from the original hash anchored on-chain. Upon executing a verification query, the smart contract compares these two strings. A hash mismatch immediately falsifies the dataset's integrity, automatically triggering contract execution logic that suppresses demand response financial rewards, emits a DataTampered event to warn network operators, and prompts a state rollback to isolate the corrupted block.

When a data batch size is large, a single point of data tampering can nave a severe impact. Since the entire batch is represented on-chain by a single cryptographic hash, tampering with a single data point invalidates the hash for the entire batch. As a result, the processing logic will reject the whole dataset, causing the aggregator to lose financial rebates or demand response credits for the entire duration of that batch, and penalizing untampered data points. To mitigate this issue without losing the cost-saving benefits of large batch sizes, a viable solution is to employ Merkle Trees instead of simple linear hashing. By structuring the energy readings as leaves in a Merkle Tree and anchoring only the Merkle Root on-chain, the system can pinpoint exactly which specific data point was altered. During a dispute, the aggregator only needs to provide a small Merkle proof (a path of hashes) to verify individual data points, allowing the smart contract to isolate only the tampered entry while successfully validating and rewarding the rest of the authentic data in the batch.

4. Experiments and Results

To evaluate the performance of our proposed solution, we conducted experiments using an open-source DR dataset presented in [8,43] and publicly available on the Smart Grid Smart City (SGSC) website [22]. The SGSC project represents one of the largest real-world commercial-scale evaluations of Smart Grid technologies, and the publicly released data include customer load profiles, electrical network properties, DR modeling and EV trials [43].

To use only data that are relevant to our analysis of DR programs, we selected the following datasets: (i) electricity use interval reading dataset, containing half-hourly interval measurements of consumption and generation sent by smart meters of users taking part in the SGSC program; (ii) peak event dataset, detailing the time and duration of several DR events issued during the SGSC trial; (iii) peak event response dataset, including the actual energy usage and the corresponding baseline load for users who agreed to the consumption reduction request.

Table 2 shows sample entries from the peak event dataset. Each event is uniquely identified by its corresponding key (i.e., an ID), and it is associated with two timestamps, indicating the event start and end time, respectively. The SGSC trial defined two categories of DR events: the tariff-based dynamic peak pricing (DPP) and the incentive-based dynamic peak rebate (DPR), as indicated in the "Event Type" column.

Table 2. Sample entries from the Peak Events dataset. Two types of DR event are defined: dynamic peak pricing (DPP) and dynamic peak rebate (DPR).

Event ID	Event Type	Start Timestamp	End Timestamp
1000241	DPR	25 January 2013 13:00	25 January 2013 17:00
1000242	DPR	8 February 2013 13:30	8 February 2013 17:30
1000243	DPP	26 February 2013 13:00	26 February 2013 15:00

Table 3 shows sample entries from the peak response dataset, where customers' identities were substituted by a numerical ID. Similarly, the DR event customers participated in is indicated by the event ID. The remaining columns show the measured consumption value during the event (in kWh) and the reference baseline, defined as the customer's energy use before the event (in kWh). Finally, the "Rebate Amount" column displays the financial incentive, in Australian dollars, which was awarded to customers who actually reduced their consumption.

Table 3. Sample entries from the Peak Response dataset. 'Actual consumption' denotes the measured consumption value during the DR event, while 'rebate amount' represents the financial remuneration provided to end users who reduced their load.

Customer ID	Event ID	Actual Consumption (kWh)	Baseline Consumption (kWh)	Rebate Amount (AUD)
10233684	1000250	1.34	5.94	20.70
10233698	1000248	1.86	5.33	15.60
10233698	1000266	2.10	7.90	26.10

Initially, to ensure the integrity of our economic analysis, we filtered the SGSC dataset to obtain only data about DPR events. In fact, DPR events in the dataset are associated with a rebate amount, whereas DPP events adopt the strategy of dynamically changing the electricity price, therefore providing an indirect financial advantage for reducing energy load. The analysis of the peak events dataset revealed that the SGSC project defined a total of 18 DPR events. However, the peak response dataset contains data for only 15 events, and among these, complete data are available for only 13 events.

Subsequently, our analysis evaluated the energy and rebate values. To determine the amount of demand reduction during the DR events, we subtracted the consumption values from the baseline values. The results showed that three DPR events exhibit a ratio of more than 85% of participants who did not reduce their demand. This unusually high ratio suggests that either the events were not properly notified to customers or there was an error in data collection and reporting. Consequently, we did not consider these events in our analysis, resulting in a total of 10 DPR events. This cleaning step aligns with the observations made in the original SGSC dataset study, which reported exceptionally high standard deviation in user responses [8]. Such variance was attributed to outliers and potential issues with baseline estimation, which is a known challenge in demand response schemes. By excluding events with irregular distribution patterns where the vast majority of users failed to respond, we ensure that our evaluation focuses on scenarios where the demand response signals were successfully delivered and acted upon, thereby providing a robust test for the certification layer.

Following this, we identified and removed all entries in which the rebate amount is positive even though the corresponding user did not actually reduce consumption. This occurrence also suggests an error in recording data, as the financial incentive is only given when the customer complies with the reduction request.

A final data cleaning procedure intended to eliminate points associated with technical faults or non-residential profiles was applied. Following the methodology established for the SGSC dataset [8], we excluded cases where the baseline estimation was lower than 0.01 kWh, as such values suggest either faulty estimations or vacant properties. Regarding high-energy outliers, the 10 kWh threshold for a single event duration was selected based on Australian residential averages. Since a typical household in Australia consumes between 20 and 25 kWh per day [44,45], a recorded consumption exceeding 10 kWh in a 3 to 4 h window is statistically likely to be an error or a non-domestic user. These thresholds ensure that the dataset remains representative of the residential energy community target.

After these pre-processing operations, we obtained a complete and clean set of data for 10 DPR events, involving 1206 unique customers and with a total of 7339 entries. Table 4 shows the statistical indices for baseline, consumption, reduction and rebate values for the cleaned peak response dataset.

Table 4. Summary statistics for the cleaned peak response dataset.

	Baseline (kWh)	Consumption (kWh)	Reduction (kWh)	Rebate (AUD)
mean	4.32	2.02	2.30	10.69
std	2.55	1.73	2.24	9.52
min	0.01	0.01	−7.66	0.00
25%	2.14	0.79	0.59	2.70
50%	3.96	1.48	1.82	8.20
75%	6.30	2.73	3.74	16.80
max	9.99	9.93	9.51	42.80

In the next phase of our analysis we determined the total number of participants for each DPR event and, based on the results, we classified the size of the set of customers into three categories: (i) small, for users up to 200 members; (ii) medium, for users between 200 and 900 members; and (iii) large, for users over 900 members. Figure 4 shows the bar plot of the number of participants for each event and the corresponding size.

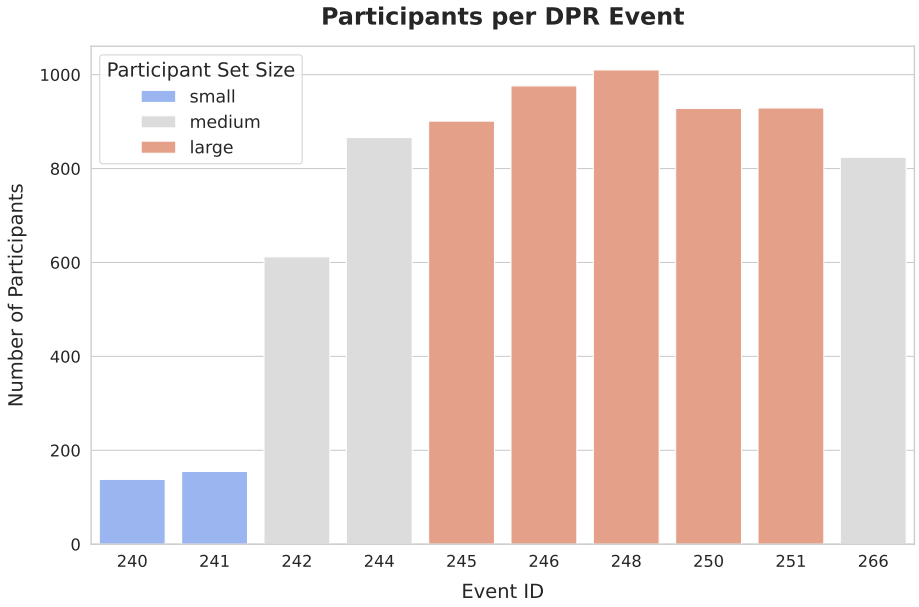

Figure 4. Number of participants for each DPR event, and the corresponding set size. Sizes were classified into three size categories: small (up to 200 members), medium (between 200 and 900 members), and large (over 900 members).

Subsequently, our analysis focused on temporal information. The ten DPR events were issued over the span of just over a year, from January 2013 to February 2014. Their total duration was three or four hours and they were predominantly called at mid-day. Finally, we observed that all events took place on weekdays (Wednesday, Thursday, and Friday), and they were more frequently called in summer and autumn. Figure 5 shows the temporal data and the time span for each DPR event.

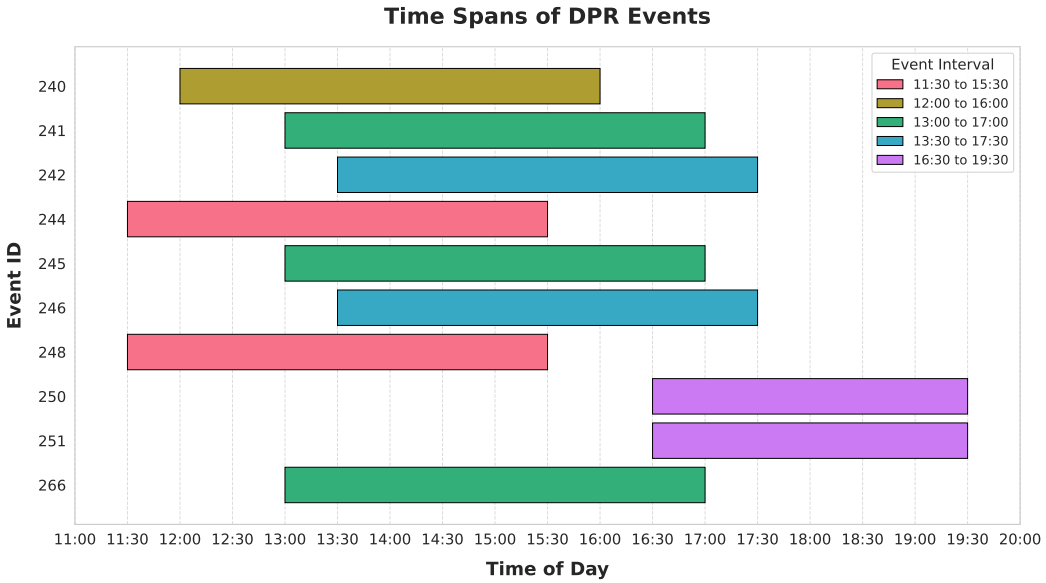

Figure 5. Time spans for the DPR events. Bars having the same color correspond to events activated during the same interval of time. The majority of DPR events were issued between 13:00 and 17:30.

With reference to the data in the cleaned dataset, timestamped information that our proposed system tracks and for which it certifies integrity includes: start and end timestamp of the DPR event, and for each community member: (i) timestamp of when they sent their reply to participate in the program, and (ii) their timestamped energy measurements for the entire duration of the event. In the SGSC trial, smart meters transmitted timestamped

energy readings at 30-min intervals. Consequently, based on the number of participants and the event duration, we determined the total amount of timestamped data to certify for each DPR event. Table 5 shows such data (the event ID was shortened to the last three digits for brevity). Additionally, the percentage of users who, according to the dataset, actually reduced their load demand is also reported. Figure 6 shows the key information from the Table 5: for each DPR event, the total amount of timestamped data to certify and the ratio of participants who actually reduced and did not reduce load demand, respectively.

Table 5. Summary of DRP event dates, durations, and the total number of participants. For each event, the percentage of users who actually reduced their consumption is also provided. The 'Timestamps' column indicates the total volume of timestamped data requiring certification.

Event ID	Date	Duration	Participants	Reducing	Timestamps
240	17 January 2013	4 h	138	74%	1244
241	25 January 2013	4 h	155	87%	1397
242	8 February 2013	4 h	612	91%	5510
244	7 March 2013	4 h	866	94%	7796
245	13 March 2013	4 h	901	95%	8111
246	22 March 2013	4 h	976	89%	8786
248	28 March 2013	4 h	1010	86%	9092
250	24 July 2013	3 h	928	93%	6498
251	2 August 2013	3 h	929	95%	6505
266	13 February 2014	4 h	824	94%	7418

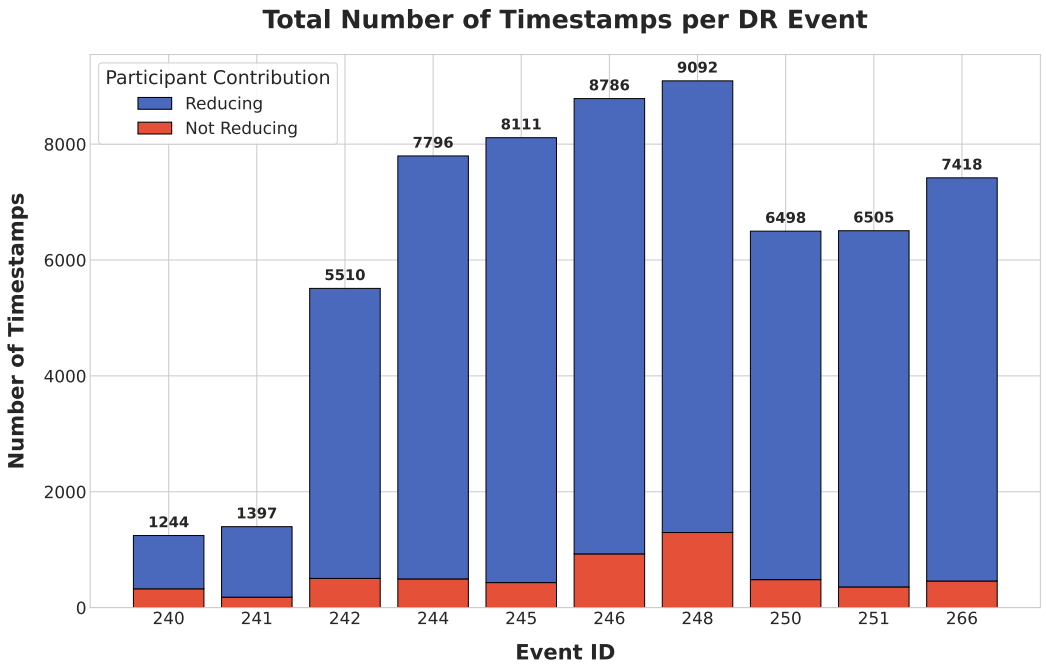

Figure 6. Stacked bar chart illustrating the total number of timestamps generated per DPR event. The blue component represents the number of timestamps from customers who reduced their consumption, while the red component represents timestamps from those who did not reduce their load.

In our simulation, the aggregator executed on a central server and activated the DPR events by sending a request to the communicator. The communicator operated on a cloud node, for resilient and reliable interaction with the community members, and transmitted users' replies back to the aggregator. Energy data readings were received and stored in the local database by the aggregator. Subsequently, based on predetermined strategies, the aggregator created energy data batches and invoked the blockchain module that was responsible for generating the data hash, storing the hash on the blockchain and returning the transaction receipt to the aggregator. Finally, the aggregator submitted the energy data, the corresponding hash and the receipt to the cloud database.

In our experiments, the data batches served as input for the SHA-256 hash function, and the resulting hash value was transmitted to the DataHashStorage smart contract for immutable storage on the blockchain. The contract was deployed on the Ethereum Sepolia testnet, and it is publicly available at the following address: 0x2a5A789beBeE29c8fc824d76a82D3aE787b6d1F5. The source code of the DataHashStorage contract and the full history of transactions can be reviewed on the corresponding Sepolia Etherscan page (https://sepolia.etherscan.io/address/0x2a5a789bebee29c8fc824d76a82d3ae787b6d1f5, accessed on 13 November 2025).

The contract implements the storeDataHash(bytes32 _dataHash) function, which takes as input the SHA-256 value and records it in an internal mapping. This mapping provides a persistent storage of key–value pairs, where each hash is associated with the timestamp of the block containing the hash writing transaction.

To assess blockchain fees and block confirmation times, we conducted a total of 1093 hash writing transactions across the Sepolia network. We requested the transactions over different weekdays, with a success rate of 100%, and evaluated the total costs associated with blockchain interactions by extracting the transaction fee from the receipt, in Gwei (1 Gwei $= 1 \times 10^{-9}$ ETH). To derive AUD fees from the corresponding Gwei amounts we used CoinGecko (https://www.coingecko.com (accessed on 3 April 2026)), a cryptocurrency data aggregator exposing a publicly available API which can be used to obtain real-time exchange rates, including ETH–AUD. The results of our experiments are reported in Figure 7, which displays the transaction fee distribution in Gwei and AUD, and Table 6, which shows the key statistical indices for block confirmation times and fees.

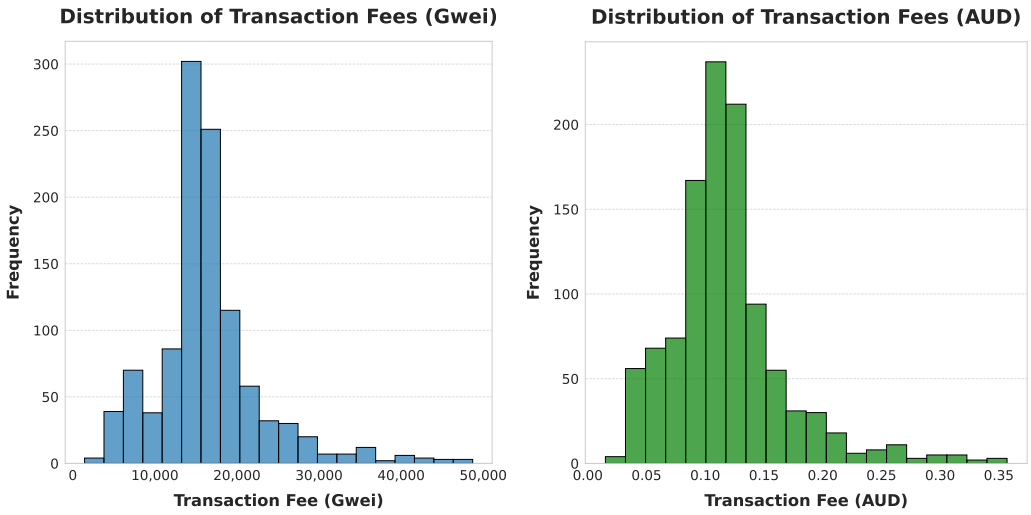

Figure 7. Distribution of blockchain transaction fees in Gwei (**left**) and Australian dollars (**right**).

Table 6. Summary statistics for blockchain transaction data.

Statistics	Transaction		
	Confirmation Time (Seconds)	Fee (Gwei)	Fee (AUD)
mean	13.38	16,305	0.120
std	5.77	6591	0.049
min	2.34	1435	0.015
25%	11.50	13,335	0.097
50%	13.19	15,662	0.116
75%	13.31	18,254	0.134
max	108.67	48,717	0.357

To optimize hash writing requests, balancing cost and performance, the aggregator used thresholds to determine the total number of entries to include in a single energy data batch. In our simulation, we considered four different thresholds: 50, 75, 125 and 175. The corresponding total transaction fees for each DPR event were derived, and are displayed in Table 7. For reference, the total and mean values for rebate amounts are also provided. Since the SGSC trial took place from 2013 to 2014, to account for inflation the reported rebate values were updated using the official Consumer price index (CPI) rates provided by Australian government portals [46,47]. Notably, the certification costs represented less than 0.1% of the total rebate amount when using higher batch thresholds (125 and 175).

The selection of data batching thresholds (125 and 175) evaluated in these experiments represents an empirical exploration of the core system trade-offs. Specifically, configuring a lower threshold decreases certification latency (the time elapsed before a data entry is anchored to the public ledger) at the expense of an increased number of transactions, which elevates cumulative gas costs during peak hours. Conversely, higher thresholds optimize economic expenditures but expose the system to prolonged periods of uncertified local buffering, making data state recovery more reliant on steady network conditions during batch initialization.

Table 7. Summary of total blockchain transaction costs by threshold. Rebate values, adjusted for inflation, are provided for reference.

ID	Rebate Amount (AUD)		Total Cost by Threshold (AUD)			
	Total	**Mean**	**50**	**75**	**125**	**175**
240	825	5.98	2.88	1.92	1.08	0.84
241	2159	13.92	3.24	2.16	1.32	0.84
242	8491	13.88	13.20	8.76	5.28	3.72
244	12,584	14.53	18.60	12.36	7.44	5.28
245	14,315	15.89	19.44	12.96	7.68	5.52
246	13,561	13.89	21.00	14.04	8.40	6.00
248	11,285	11.17	21.72	14.52	8.64	6.12
250	14,945	16.10	15.48	10.32	6.12	4.44
251	16,682	17.96	15.60	10.32	6.24	4.44
266	12,840	15.58	17.76	11.76	7.08	5.04

To evaluate the performance of the dispute resolution workflow, a series of simulation experiments were conducted to benchmark the total verification latency. Data retrieval from the remote cloud server was tested using an enterprise-grade storage environment, yielding response times between 100 and 200 ms. To quantify the blockchain reading overhead, the getBlockTimestampFromHash() function was invoked across 10,000 test iterations on the Ethereum Sepolia network. The results indicated that the reading request on the distributed ledger is finalized within 150 to 400 ms. Finally, the computational overhead for re-calculating the SHA-256 hash of the energy data batch was measured as sub-millisecond, resulting in a total verification delay consistently lower than 600 ms.

Although the experiments were executed on the Sepolia testnet, the observed variance in transaction fees reflects real-world blockchain network congestion. When public network traffic spikes, the aggregator face a critical trade-off: it must either pay higher gas premiums to maintain a low commit latency, or queue transactions using the internal buffering mechanism until network conditions normalize. This buffering introduces a temporary latency penalty for data certification but successfully prevents cost overruns, validating the asynchronous robustness of the framework.

5. Discussion

A total of 1093 hash writing transactions were requested and finalized on the Ethereum Sepolia blockchain. To reproduce the same temporal conditions of the SGSC trial, transactions were executed during weekdays and over the same time interval, spanning from 11:30 to 19:30 in the Australian timezone (see Figure 5). Although the SGSC trial was in effect in the years 2013–2014, it is still a realistic assumption that DR events are requested in similar time intervals, when energy consumption is typically higher. Moreover, during midday the energy production from solar photovoltaic panels reaches its peak, which is significant especially in summer, when the majority of DPR events were requested. A high availability of energy output from renewable energy sources ensures more options for community members willing to participate in the DR program, especially prosumers equipped with ESS and EV charging stations (in the V2G scenario).

Table 4 reveals that, on average, customers achieved a reduction of 2.30 kWh, with 75% of all entries reporting a value as high as 3.74 kWh. Since the average baseline value, i.e., the estimation of user consumption without DPR events, is 4.32 kWh, the results in the table confirm that the DPR request and associated rewards had a significant impact on modifying energy consumption patterns and ensuring that grid congestion is mitigated. While the presence of negative reduction values signifies that some users increased their demand load, our analysis in the cleaned datasets demonstrated that the vast majority of customers who agreed to take part in the DR program actually reduced consumption during the DPR event (see Table 5). Specifically, the average ratio of complying users was 89.8%, with percentage exceeding 90% for 6 out of 10 events. Such data represent an important confirmation that DR programs are successful and that financial rewards make an important contribution to ensuring that users respond to the request of modifying their energy usage patterns. It is important to note that our approach continues to track and certify energy data for all users who agreed to the DR request, even when some of them do not actually reduce consumption during the DR event. This guarantees the system can prove that non-complying users are not entitled to a reward, should a dispute arise.

The statistical analysis of block confirmation times and transaction fees for the 1093 hash writing operations provides a solid basis for evaluating performance and economic feasibility in a real-world scenario. The experimental results show that the mean time to finalize a block is 13.38 s, with 75% of all transactions taking at most 13.31 s (see Table 6). While we observed the presence of some outliers, with one hash writing operation taking approximately 109 s to complete, these confirmation delays have no impact on the incentive settlement process. The validity of a user's participation in a Demand Response event is determined by the original smart meter timestamp, which is generated at the moment of energy measurement. Because the certification process is decoupled from the actual data collection, a delay in recording the hash on the blockchain does not cause users to miss event deadlines. The aggregator acts as a secure buffer, ensuring that all data is preserved locally until the asynchronous blockchain transaction is finalized. Additionally, whenever transactions are not confirmed due to factors such as network faults, the aggregator resubmits the hash writing request to the blockchain module for the same data batch. Finally, when the transaction is completed and the corresponding receipt is returned, the data batch is linked to an immutable record that certifies energy data integrity. This property is essential for verifiability and for complying with legal requirements if disputes arise. In fact, any malicious or unauthorized modification to the cloud-based data can be detected, because the resulting hash would differ from the one retrieved from the blockchain. Conversely, identical hashes confirm data integrity, and enable dispute resolution.

Costs for executing transactions on the Ethereum blockchain are paid in ETH. A transaction fee distribution originates from the same hash writing operation because the fee

for using the blockchain depends on the state of the network, resulting in higher prices during periods of congestion. The difference between the Gwei and corresponding AUD distributions (see Figure 7) is due to the fact that ETH/AUD exchange rates are continuously subject to fluctuations. In our simulations, we observed that the average cost for finalizing a blockchain writing operation is AUD 0.12, and 75% of transactions' costs were less than AUD 0.135 (see Table 6). Moreover, the lowest and highest recorded costs were AUD 0.357 and AUD 0.015, respectively. We used the mean fee value to estimate the total cost needed to provide integrity certification for energy data in every DPR event, and the results are shown in Table 7, divided by batch size threshold. The table analysis reveals that the total cost for providing proof of integrity for the entire set of energy data during a DPR event is lower than the average rebate amount, adjusted for inflation, in the vast majority of the scenarios. Specifically, the total costs remain consistently lower with the 125 and 175 thresholds, and exceed the average in only two and five occurrences for the 75 and 50 thresholds, respectively. Overall, total cost is less than the average rebate in 80% of the cases, and even when the cost is higher than the average, it represents only a small fraction of the total rebate amount awarded to participants (see Table 7). Certification cost-to-total rebate ratios range from 0.03% to 0.35%, and remain consistently lower than 0.10% for the 125 and 175 thresholds.

Higher thresholds correspond to larger data batches and less frequent hash writing requests, resulting in lower overall costs. However, this introduces a critical latency trade-off: larger thresholds increase the accumulation time before a hash is anchored on-chain, thereby delaying finality and lengthening the time window wherein uncertified data remains vulnerable in the aggregator's local storage. The threshold choice must balance this certification latency against operational cost savings and cloud retrieval data overhead during dispute resolution, adjusting dynamically based on system requirements. To address this challenge, our approach provides a high level of flexibility: depending on the number of participants (which can significantly vary from event to event, see Figure 4) and blockchain costs, the thresholds can be dynamically changed to adapt to different conditions and achieve an optimal balance between cost and performance.

The experimental results regarding verification latency demonstrate the practical feasibility of our hybrid storage model for rapid dispute resolution. With a total delay from dispute initiation to verification completion of under 600 ms, the system provides a near-instantaneous audit trail for Energy Community members. By combining high-speed cloud retrieval with public blockchain immutability, our model ensures that data integrity can be verified with minimal operational overhead, thereby fostering trust without compromising system responsiveness.

The proposed architecture functions as a resilient, asynchronous integrity layer for energy data. The aggregator acts as a resilient buffer, retaining energy readings in a local database before initiating the batching and hashing process. This design ensures that even during periods of network congestion or temporary communication faults, data loss is prevented. Nevertheless, extended network degradation can cause a backlog of unhashed records, inflating local database overhead and exponentially increasing the subsequent cost and latency required to clear the transaction pool when network connectivity stabilizes. Such a mechanism is highly effective for demand response programs where financial rebates are settled retroactively, rather than requiring immediate real-time on-chain finality, provided the aggregator's local storage capacity can withstand prolonged communication dropouts. By committing only a 32-byte hash for each data batch, the on-chain footprint is minimized, enabling the system to scale efficiently with the number of participants and the granularity of energy readings. Although the experiments were conducted on the Sepolia testnet to validate the architectural logic, the inclusion of 1093 transactions across

different weekdays provides a representative assessment of the system's reliability and the aggregator's ability to handle variable network conditions.

Comparative Evaluation of Architectural Alternatives

To further validate the economic and technical feasibility of our framework, we compare our dual-storage architecture against three alternative baselines frequently discussed in blockchain literature: Full On-Chain storage, Private Sidechains, and Decentralized Storage (IPFS/Filecoin). Table 8 provides a multi-dimensional analysis based on reported performance metrics and theoretical cost projections.

Table 8. Multi-dimensional Comparison of Storage Strategies.

Dimension	Full On-Chain	Private Sidechain	IPFS + Filecoin	Proposed (Hybrid)
Estimated Cost	Prohibitive (>10^5)	Low (OpEx only)	Moderate (Fees/Deals)	**Minimal**
Security/Trust	Maximum (Public)	Low (Operator-led)	High (Decentralized)	**High (Public)**
Retrieval Latency	Very High	Very Low	High	**Very Low**
Storage Overhead	Extreme (Every node)	High	Moderate	**Low (Cloud)**
Verification	On-chain Consensus	Trusted Authority	Proof-of-Spacetime	**Merkle Proof**

As illustrated in Table 8, while storing full datasets directly on-chain offers maximum redundancy, it is economically impractical for high-frequency energy data. Recent studies confirm that storing raw data on public blockchains like Ethereum leads to prohibitive costs as the total ledger size continues to grow, exceeding 1.4 TB by the end of 2025 [48]. Although private sidechains offer lower operational costs, they typically involve governance by a central operator, which introduces trust issues similar to traditional databases [49]; if the DR aggregator controls the sidechain, they retain the power to manipulate records, defeating the purpose of immutable third-party certification.

Furthermore, while decentralized storage solutions like IPFS and Filecoin have been proposed, they introduce significant retrieval latency—often exceeding 200 ms compared to the millisecond-scale response of cloud providers [50]. These platforms also struggle with long-term data persistence for time-sensitive utility data, as a significant portion of peer-to-peer nodes remains online for less than a week without complex incentive management [51]. By utilizing a hybrid model that pairs efficient cloud storage with public blockchain hashes, our approach achieves the sub-second retrieval latency required for rapid dispute resolution while maintaining the high-integrity guarantees of a public ledger.

To ensure cost-effectiveness in practical deployments, the proposed architecture enables high flexibility by allowing thresholds to be dynamically changed based on the number of participants and transaction costs. Since the size of the energy community is known a priori, the aggregator can calculate optimal batch sizes before a Demand Response event begins. Additionally, to mitigate gas fee volatility, a polling mechanism can be employed to monitor the current blockchain status and ensure transactions are finalized during periods of lower network activity. This creates an inherent trade-off between operational cost and dispatch latency: delaying transaction finalization to avoid gas spikes directly prolongs the uncertified status of the data batch. If network congestion persists indefinitely, the system must choose between absorbing elevated transaction costs or violating the maximum allowed verification latency bounds required by regulatory frameworks. Such a strategy is particularly effective when dealing with the half-hourly readings found in the SGSC trial, but it is equally applicable to energy systems where measurements are transmitted using significantly shorter time intervals, e.g., every five minutes.

Overall, our approach of storing on-chain only 32-byte hashes rather than the complete record of data achieves a significant reduction in blockchain utilization costs while preserv-

ing data confidentiality. The empirical results demonstrated that transaction fees remain comparatively low with respect to DPR events reward amounts, despite cost fluctuations due to blockchain congestion. Additionally, hashes are stored on a public blockchain, enabling independent data verifiability by authorized entities when disputes arise. This mechanism promotes trust among all EC members and represents a key advantage in contrast to third-party and commercial certification systems, which can define arbitrary costs and terms of use for their services and therefore cannot be fully trusted.

Beyond cost reduction, the proposed architecture demonstrates favourable scalability and energy efficiency properties. By aggregating multiple measurements into a single blockchain transaction, the system significantly reduces the number of required on-chain operations, enabling it to handle high-frequency data streams without overwhelming the network. This batching strategy ensures that the system can scale with the number of participants and measurement granularity, which are expected to grow in future energy communities.

From an energy efficiency perspective, reducing the number of blockchain transactions directly lowers the computational effort required by the underlying consensus mechanism, thereby decreasing the overall energy footprint of the system. This aspect is particularly important given the increasing scrutiny of blockchain sustainability. The proposed hybrid approach therefore strikes a balance between the strong integrity guarantees of blockchain and the need for lightweight, sustainable operation in large-scale energy systems.

From a privacy perspective, the public visibility of the blockchain does not compromise user confidentiality, as only 32-byte hashes are recorded on the distributed ledger. These hashes represent aggregated batches of data, which prevent the extraction of individual consumption profiles. While the cloud storage maintains the complete set of raw data and metadata, it is protected by robust authorization and access control protocols. Compliance with data protection regulations, is managed at the organizational level by the Energy Community, which ensures that the aggregator collects and processes data according to established legal frameworks for consumer protection.

In our experiments, we used publicly available datasets officially released by the Australian government, which correspond to a real-world implementation of an incentive-based DR program, involving thousands of households over the course of more than a year. Consequently, our cost analysis can be effectively used as a basis to obtain realistic cost estimates to certify data integrity in DR programs. Additionally, these estimations enable cost projection for different DR configurations, with varying number of participants, duration and frequency of DR events. This informed analysis represents a key aspect for optimal investment planning and the definition of DR program parameters.

Results analysis shows that our proposed approach achieves feasible operational costs even in the presence of high variability for blockchain transaction fees. The strategy of aggregating energy data in batches and committing on the blockchain only the corresponding hash enables the proof of integrity for the data at a minimal cost. Finally, the definition of thresholds for batch sizes enables the optimization of writing requests, which avoids blockchain congestion and provides a high degree of flexibility to system operators.

6. Conclusions

This article presented a framework that leverages blockchain properties to provide proof of integrity for data in energy communities participating in incentive-based demand response programs. The designed system achieves data certification using the decentralized blockchain network, tackling a significant limitation of centralized approaches, which require trust in an external authority. The proposed system requests on-chain hash writing transactions of energy data in the context of demand response events. The hashes recorded

on the blockchain represent an immutable record that can be used to certify the validity of energy data and detect malicious data alteration in the database.

Experimental assessment of the proposed system was conducted on a real-world energy dataset released by government agencies that designed and implemented demand response trials spanning over a year. Results show that our approach achieves efficient and cost-effective data management. The definition of configurable thresholds enables a high degree of flexibility, allowing operators to dynamically balance the tripartite trade-off between verification latency, network congestion constraints, and blockchain transaction costs. This multi-dimensional adjustability allows for effectively managing fluctuations in both the number of participants and fluctuating network conditions.

Crucially, the insights derived from our study pave the way for broader applicability of blockchain-based data certification in distributed systems. Future work should investigate the edge-to-aggregator communication reliability under volatile network conditions, exploring how cellular or mesh network latency affects real-world batching intervals. Ultimately, in the context of distributed data systems, having a reliable and cost-effective strategy for proving data integrity and obtaining an independently verifiable audit trail is paramount, and it promotes trust among stakeholders without relying on a central authority.

Author Contributions: Conceptualization, E.T.; methodology, D.M., A.M. and E.T.; software, D.M.; validation, D.M., A.M. and E.T.; writing—original draft preparation, D.M. and A.M.; writing—review and editing, E.T.; visualization, D.M. and A.M.; supervision, E.T.; project administration, E.T.; funding acquisition, E.T. All authors have read and agreed to the published version of the manuscript.

Funding: This research received no external funding.

Data Availability Statement: The original data utilized in this study are openly available from the Australian Government data portal at https://data.gov.au/data/dataset/smart-grid-smart-city-customer-trial-data (accessed on 13 November 2025).

Acknowledgments: The authors acknowledge the support of the University of Catania PIACERI Project "TEAMS".

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

References

1. Slameršak, A.; Kallis, G.; O'Neill, D.W. Energy requirements and carbon emissions for a low-carbon energy transition. *Nat. Commun.* **2022**, *13*, 6932. [CrossRef] [PubMed]
2. Bartie, N.; Cobos-Becerra, Y.; Fröhling, M.; Schlatmann, R.; Reuter, M. The resources, exergetic and environmental footprint of the silicon photovoltaic circular economy: Assessment and opportunities. *Resour. Conserv. Recycl.* **2021**, *169*, 105516. [CrossRef]
3. (IRENA). Renewable Energy Statistics. 2023. Available online: https://www.irena.org/Publications/2023/Jul/Renewable-energy-statistics-2023 (accessed on 13 November 2025).
4. International Energy Agency (IEA). Renewables 2024. 2024. Available online: https://www.iea.org/reports/renewables-2024 (accessed on 13 November 2025).
5. Marletta, D.; Midolo, A.; Tramontana, E. Detecting Photovoltaic Panels in Aerial Images by Means of Characterising Colours. *Technologies* **2023**, *11*, 174. [CrossRef]
6. Zhang, C.; Zhang, Y.J. Optimal solar panel placement in microgrids. In Proceedings of the IEEE International Conference on Smart Grid Communications (SmartGridComm), Sydney, Australia, 6–9 November 2016; pp. 376–381. [CrossRef]
7. Wang, W.; Yuan, B.; Sun, Q.; Wennersten, R. Application of energy storage in integrated energy systems—A solution to fluctuation and uncertainty of renewable energy. *J. Energy Storage* **2022**, *52*, 104812. [CrossRef]
8. Antonopoulos, I.; Robu, V.; Couraud, B.; Flynn, D. Data-driven modelling of energy demand response behaviour based on a large-scale residential trial. *Energy AI* **2021**, *4*, 100071. [CrossRef]
9. Honarmand, M.E.; Hosseinnezhad, V.; Hayes, B.; Shafie-Khah, M.; Siano, P. An overview of demand response: From its origins to the smart energy community. *IEEE Access* **2021**, *9*, 96851–96876. [CrossRef]
10. Muthirayan, D.; Baeyens, E.; Chakraborty, P.; Poolla, K.; Khargonekar, P.P. A minimal incentive-based demand response program with self reported baseline mechanism. *IEEE Trans. Smart Grid* **2019**, *11*, 2195–2207. [CrossRef]

11. Meng, L.; Chen, L. Analysis of Client-Side Security for Long-Term Time-Stamping Services. In *Proceedings of the Applied Cryptography and Network Security*; Sako, K., Tippenhauer, N.O., Eds.; Springer: Cham, Switzerland, 2021; pp. 28–49. [CrossRef]
12. Wilson, K.S. Blockchain with Daisy Chained Records, Document Corral, Quarantine, Message Timestamping, and Self-Addressing. US Patent 11,444,776, 13 September 2022.
13. Marletta, D.; Midolo, A.; Tramontana, E. A Blockchain-Based Strategy for Certifying Timestamps in a Distributed Healthcare Emergency Response Systems. *Future Internet* **2025**, *17*, 210. [CrossRef]
14. Zhang, S.; Zhou, Y.; Wang, Z.; Huang, Y.; Zhang, C.; Wei, L. An Efficient Blockchain-Based Time-Stamping Scheme Using Commitment Signatures. In Proceedings of the IEEE Global Communications Conference (GLOBECOM), Rio de Janeiro, Brazil, 4–8 December 2022; pp. 2591–2596. [CrossRef]
15. Koukaras, P.; Afentoulis, K.D.; Gkaidatzis, P.A.; Mystakidis, A.; Ioannidis, D.; Vagropoulos, S.I.; Tjortjis, C. Integrating Blockchain in Smart Grids for Enhanced Demand Response: Challenges, Strategies, and Future Directions. *Energies* **2024**, *17*, 1007. [CrossRef]
16. Pop, C.; Cioara, T.; Antal, M.; Anghel, I.; Salomie, I.; Bertoncini, M. Blockchain based decentralized management of demand response programs in smart energy grids. *Sensors* **2018**, *18*, 162. [CrossRef] [PubMed]
17. Mao, X.; Xue, M.; Zhang, T.; Tan, W.; Zhang, Z.; Pan, Y.; Wu, H.; Lin, Z. Centralized bidding mechanism of demand response based on blockchain. *Energy Rep.* **2022**, *8*, 111–117. [CrossRef]
18. Jayabalan, J.; Jeyanthi, N. Scalable blockchain model using off-chain IPFS storage for healthcare data security and privacy. *J. Parallel Distrib. Comput.* **2022**, *164*, 152–167. [CrossRef]
19. Hepp, T.; Sharinghousen, M.; Ehret, P.; Schoenhals, A.; Gipp, B. On-chain vs. off-chain storage for supply- and blockchain integration. *It-Inf. Technol.* **2018**, *60*, 283–291. [CrossRef]
20. Xu, C.; Zhang, C.; Xu, J.; Pei, J. SlimChain: Scaling blockchain transactions through off-chain storage and parallel processing. *Proc. VLDB Endow.* **2021**, *14*, 2314–2326. [CrossRef]
21. Iansiti, M.; Lakhani, K.R. The truth about blockchain. *Harv. Bus. Rev.* **2017**, *95*, 118–127.
22. Australian Government—Department of Climate Change, Energy, the Environment and Water. Smart-Grid Smart-City Customer Trial Data. 2022. Available online: https://data.gov.au/data/dataset/smart-grid-smart-city-customer-trial-data (accessed on 13 November 2025).
23. Bayer, D.; Haber, S.; Stornetta, W.S. Improving the Efficiency and Reliability of Digital Time-Stamping. In *Proceedings of the Sequences II*; Capocelli, R., De Santis, A., Vaccaro, U., Eds.; Springer: New York, NY, USA, 1993; pp. 329–334. [CrossRef]
24. Meng, L.; Chen, L. Long-Term Time-Stamping Schemes Based on MACs, Archives, and Transient Keys. In Proceedings of the IEEE International Conferences on Internet of Things (iThings) and IEEE Green Computing and Communications (GreenCom) and IEEE Cyber, Physical and Social Computing (CPSCom) and IEEE Smart Data (SmartData) and IEEE Congress on Cybermatics, Copenhagen, Denmark, 19–22 August 2024; pp. 535–542. [CrossRef]
25. Vigil, M.; Weinert, C.; Demirel, D.; Buchmann, J. An efficient time-stamping solution for long-term digital archiving. In Proceedings of the IEEE International Performance Computing and Communications Conference (IPCCC), Austin, TX, USA, 5–7 December 2014; pp. 1–8. [CrossRef]
26. Meng, L.; Chen, L. A Blockchain-Based Long-Term Time-Stamping Scheme. In *Proceedings of the Computer Security—ESORICS*; Atluri, V., Di Pietro, R., Jensen, C.D., Meng, W., Eds.; Springer: Cham, Switzerland, 2022; pp. 3–24. [CrossRef]
27. Agarwal, N.; Wongthongtham, P.; Khairwal, N.; Coutinho, K. Blockchain Application to Financial Market Clearing and Settlement Systems. *J. Risk Financ. Manag.* **2023**, *16*, 452. [CrossRef]
28. Zhang, Y.; Xu, C.; Cheng, N.; Li, H.; Yang, H.; Shen, X. Chronos: An Accurate Blockchain-Based Time-Stamping Scheme for Cloud Storage. *IEEE Trans. Serv. Comput.* **2020**, *13*, 216–229. [CrossRef]
29. Qian, B.; Luo, Y.; Ou, J.; Xiao, Y.; Hu, H. IoETTS: A Decentralized Blockchain-based Trusted Time-stamping Scheme for Internet of Energy. *J. Internet Technol.* **2023**, *24*, 519–529. [CrossRef]
30. Bâra, A.; Oprea, S.V. Enabling coordination in energy communities: A Digital Twin model. *Energy Policy* **2024**, *184*, 113910. [CrossRef]
31. Lucas, A.; Geneiatakis, D.; Soupionis, Y.; Nai-Fovino, I.; Kotsakis, E. Blockchain Technology Applied to Energy Demand Response Service Tracking and Data Sharing. *Energies* **2021**, *14*, 1881. [CrossRef]
32. Dey, B.; Misra, S.; Garcia Marquez, F.P. Microgrid system energy management with demand response program for clean and economical operation. *Appl. Energy* **2023**, *334*, 120717. [CrossRef]
33. Harsh, P.; Das, D. Optimal coordination strategy of demand response and electric vehicle aggregators for the energy management of reconfigured grid-connected microgrid. *Renew. Sustain. Energy Rev.* **2022**, *160*, 112251. [CrossRef]
34. Lu, X.; Li, H.; Zhou, K.; Yang, S. Optimal load dispatch of energy hub considering uncertainties of renewable energy and demand response. *Energy* **2023**, *262*, 125564. [CrossRef]
35. Gupta, S.; Maulik, A.; Das, D.; Singh, A. Coordinated stochastic optimal energy management of grid-connected microgrids considering demand response, plug-in hybrid electric vehicles, and smart transformers. *Renew. Sustain. Energy Rev.* **2022**, *155*, 111861. [CrossRef]

36. Jiang, W.; Wang, X.; Huang, H.; Zhang, D.; Ghadimi, N. Optimal economic scheduling of microgrids considering renewable energy sources based on energy hub model using demand response and improved water wave optimization algorithm. *J. Energy Storage* **2022**, *55*, 105311. [CrossRef]
37. Zhang, H.; Qiu, D.; Kok, K.; Paterakis, N.G. Reliability assessment of multi-agent reinforcement learning algorithms for hybrid local electricity market simulation. *Appl. Energy* **2025**, *389*, 125789. [CrossRef]
38. European Parliament. Directive (EU) 2018/2001 of the European Parliament and of the Council of 11 December 2018 on the Promotion of the Use of Energy from Renewable Sources. Available online: https://eur-lex.europa.eu/eli/dir/2018/2001/oj (accessed on 13 November 2025).
39. Leal Filho, W.; Trevisan, L.V.; Salvia, A.L.; Mazutti, J.; Dibbern, T.; de Maya, S.R.; Bernal, E.F.; Eustachio, J.H.P.P.; Sharifi, A.; Alarcón-del Amo, M.d.C.; et al. Prosumers and sustainable development: An international assessment in the field of renewable energy. *Sustain. Futur.* **2024**, *7*, 100158. [CrossRef]
40. López, I.; Goitia-Zabaleta, N.; Milo, A.; Gómez-Cornejo, J.; Aranzabal, I.; Gaztanaga, H.; Fernandez, E. European energy communities: Characteristics, trends, business models and legal framework. *Renew. Sustain. Energy Rev.* **2024**, *197*, 114403. [CrossRef]
41. Kerscher, S.; Arboleya, P. The key role of aggregators in the energy transition under the latest European regulatory framework. *Int. J. Electr. Power Energy Syst.* **2022**, *134*, 107361. [CrossRef]
42. Antonopoulos, A.M.; Wood, G. *Mastering Ethereum: Building Smart Contracts and Dapps*; O'reilly Media: Santa Rosa, CA, USA, 2018.
43. Motlagh, O.; Foliente, G.; Grozev, G. Knowledge-mining the Australian smart grid smart city data: A statistical-neural approach to demand–response analysis. In *Planning Support Systems and Smart Cities*; Springer: Cham, Switzerland, 2015; pp. 189–207. [CrossRef]
44. Fan, H.; MacGill, I.; Sproul, A.B. Statistical analysis of driving factors of residential energy demand in the greater Sydney region, Australia. *Energy Build.* **2015**, *105*, 9–25. [CrossRef]
45. Roberts, M.B.; Haghdadi, N.; Bruce, A.; MacGill, I. Characterisation of Australian apartment electricity demand and its implications for low-carbon cities. *Energy* **2019**, *180*, 242–257. [CrossRef]
46. Australian Government—Australian Taxation Office. Consumer Price Index (CPI) Rates. Available online: https://www.ato.gov.au/tax-rates-and-codes/consumer-price-index (accessed on 13 November 2025).
47. Australian Bureau of Statistics. Consumer Price Index, Australia. Available online: https://www.abs.gov.au/statistics/economy/price-indexes-and-inflation/consumer-price-index-australia (accessed on 13 November 2025).
48. Zhu, T.; Huang, Y.; Liang, Z.; Qin, M.; Niu, R.; Ma, Y.; Feng, Q. Research on an On-Chain and Off-Chain Collaborative Storage Method Based on Blockchain and IPFS. *Future Internet* **2026**, *18*, 92. [CrossRef]
49. Strehle, E. Public versus private blockchains. In *BRL Working Paper*; Blockchain Research Lab: Hamburg, Germany, 2020.
50. Mandinyenya, G.; Malele, V. Comparative security and performance evaluation of ipfs and filecoin for off-chain blockchain storage. *Indones. J. Comput. Sci.* **2025**, *14*, 6241–6258. [CrossRef]
51. Wu, Z.; Lin, B.; Nikiforakis, N.; Balasubramanian, A. Degrees of Decentralized Freedom: Comparing Modern Decentralized Storage Platforms. In Proceedings of the 2025 9th Network Traffic Measurement and Analysis Conference (TMA), Copenhagen, Denmark, 10–13 June 2025; pp. 1–11. [CrossRef]