energies

MDPI

Article

Federated Distributed Scheduling for Hydrogen Production Under Renewable Variability: A Safety-Constrained Evaluation of FedAvg, FedProx, Gossip, and Local Control

Shaymaa W. Al-Shammari [1] and Moahaimen Talib [2,*]

[1] Department of Computer Engineering, College of Engineering, Al-Nahrain University, Baghdad 10072, Iraq; shaymaa.al-shammari@nahrainuniv.edu.iq
[2] College of Engineering, Al-Bayan University, Baghdad 10001, Iraq
* Correspondence: moahaimen.t@albayan.edu.iq

Abstract

Distributed hydrogen refueling stations enable the coupling of renewable generation, storage, and demand fulfillment; however, their performance depends on coordinated control under strict physical safety limits. Centralized controllers are often impractical due to privacy constraints and unreliable communication links, while unconstrained learning can reduce operating costs at the expense of unsafe pressure excursions. Therefore, this study evaluates safety-constrained coordination across multiple stations using federated learning-based distributed scheduling and benchmarks a non-federated Local Control baseline (local-only, no coordination). Using a feasibility-first rule with an acceptance threshold of $\tau = 0.2$ on the pressure violation metric ($V_p \leq 0.2$), the best feasible overall controller (Local Control) achieved a cost of 2131.83 with pressure violation $V_p = 0.172$, representing a 37.22% reduction relative to a centralized reference cost of 3396.25. Federated training with Federated Averaging and a solar–wind mixing scheme produced the best feasible federated policy (cost 2423.72, $V_p = 0.163$) with 866,688 transmitted bytes. Extensive simulations report cost, unmet demand, safety violations, and communication overhead, demonstrating that feasibility-first selection is essential because lower-cost policies can be unsafe (e.g., cost 1952.27 with $V_p = 2.63$).

Keywords: distributed scheduling; federated learning; constrained optimization; renewable energy integration; decentralized coordination

Check for updates

Academic Editor: Alberto Pettinau

Received: 21 January 2026
Revised: 14 February 2026
Accepted: 4 March 2026
Published: 11 March 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) license.

1. Introduction

Hydrogen is increasingly viewed as a practical pathway for decarbonizing sectors where direct electrification is difficult, and it is now central to many net-zero roadmaps and infrastructure plans worldwide [1]. However, scaling hydrogen supply to meet mobility and urban demand remains constrained by the operating cost of hydrogen refueling stations (HRSs) [2]. In addition, renewable generation variability and strict station-level safety limits complicate operational scheduling and control [3]. Recent work has further incorporated demand uncertainty and electricity price variation into multi-station management formulations [4]. In particular, high-pressure operation introduces safety and reliability risks that must be handled explicitly rather than treated as minor penalties during optimization [5].

As the number of stations increases, the classical assumption of a fully centralized optimizer becomes less realistic. Stations often operate under heterogeneous renewable

profiles and demand patterns, and operators are often unable or unwilling to share raw operational data due to privacy constraints and bandwidth bottlenecks. In energy applications, learning-based load forecasting from smart meter data has illustrated how computation can be performed under data locality constraints without centralizing raw measurements [6]. Building on this, federated learning (FL) has emerged as a practical privacy-preserving collaboration framework for distributed settings [7]. For HRS operation specifically, learning-based real-time station energy management has been explored under uncertainty, showing that data-driven policies can improve operational performance when exogenous signals fluctuate [8]. Nevertheless, decentralized gossip-style learning provides a serverless alternative to server-based FL but introduces topology- and communication-dependent behavior [9].

Practical deployments must also assume imperfect participation and unreliable links, which motivates aggregation rules designed to tolerate outliers or faulty updates [10]. Variance-reduced robust FL has been proposed to further stabilize training under non-i.i.d. data and adversarial behavior [11]. Robust aggregation rules have also been formalized with explicit resilience objectives in federated optimization [12]. Communication efficiency remains a dominant constraint in edge deployments, motivating mechanisms that reduce bandwidth while preserving learning utility [13]. At the system level, coordination across multiple HRSs has also been studied through multi-agent reinforcement learning formulations [14]. Event-triggered communication further supports communication-aware operation by transmitting only when updates change meaningfully [15]. Accordingly, we apply a feasibility-first accept-then-rank rule based on a prescribed pressure threshold τ.

Specifically, we enforce a feasibility-first accept-then-rank decision rule. Equations (5) and (6) screen methods by requiring $V_p(m) \leq \tau$ and then rank only the feasible candidates by operating cost to select the deployed method. This accept-then-rank principle turns evaluation from "best average cost" into an engineering decision rule aligned with safety-critical operations.

1.1. Literature Review

A large body of work targets the economic operation of HRSs as hydrogen deployment accelerates under national roadmaps and decarbonization strategies [1]. Foundational modeling and operation-strategy studies for on-site HRSs provide structured formulations for planning, dispatch, and control decisions under realistic station constraints [2]. Recent contributions further improve HRS operating economics through optimized scheduling under electricity price variation and renewable intermittency [3]. For instance, optimized operational planning for an on-site electrolytic HRS can significantly reduce daily operating costs by coordinating hydrogen production with demand estimation and electricity pricing [4]. In parallel, hydrogen infrastructure safety has been examined using quantitative risk assessment, emphasizing that high-pressure hydrogen systems require explicit mitigation actions and risk-aware design considerations [5]. While these studies strengthen station-level operation, they do not provide an operational decision framework for selecting distributed control methods under safety constraints.

On the distributed learning side, FL has matured into a widely used paradigm for privacy-preserving collaboration across edge devices, with surveys detailing mechanisms and system challenges [7]. Learning-based approaches have also been explored for real-time station energy management under uncertainty [8]. Decentralized training can deliver strong predictive performance without collecting raw local data centrally [6]. However, practical FL must confront unreliable communication and adversarial behavior. Robust aggregation methods have been proposed to tolerate outliers and poisoning, including

Byzantine-resilient defenses [10]. Communication efficiency remains a primary constraint in edge deployments, motivating compression mechanisms that reduce bandwidth [13].

Beyond station-level optimization, coordination across multiple HRSs has been studied through distributed scheduling and multi-agent reinforcement learning, where networked stations cooperate to supply fuel and provide grid services [14]. Although such coordination models capture system-level benefits, they are rarely evaluated under deployment realities (partial participation, packet dropouts, and explicit communication load accounting over time).

More broadly, event-triggered learning/control is studied as a general framework for reducing unnecessary transmissions in networked systems [15]. Related integrated energy station formulations also investigate coordinated electricity–hydrogen energy management across multiple charging/hydrogen stations [4]. Meanwhile, energy-efficient FL variants based on event-triggered communication further highlight how triggering policies can reduce communication while preserving learning behavior in distributed settings [15].

1.2. Research Gaps and Contributions

Existing studies on hydrogen refueling station (HRS) operation primarily emphasize cost-driven scheduling under operational constraints [2]. Multi-station coordination has also been investigated in electricity–hydrogen coupled settings, but the evaluation has typically been cost- or service-driven rather than framed as an explicit feasibility-first decision rule tied to high-pressure safety acceptance [4]. Similarly, coordination via distributed scheduling and multi-agent learning has been explored at the network level, yet the selection logic has rarely been expressed as an accept/reject gate that screens unsafe solutions before ranking performance [14].

The safety-oriented literature clearly establishes that hydrogen infrastructure requires explicit mitigation and risk-aware design, but it is commonly evaluated through quantitative risk metrics rather than used as an operational selection rule for control and coordination strategies [5]. In parallel, learning-based control has been applied to real-time HRS energy management under uncertainty, demonstrating that data-driven policies can improve operational outcomes when exogenous signals fluctuate [8]. However, these strands are often studied in isolation from deployment realities (imperfect participation and explicit communication accounting), making it difficult to compare coordination paradigms under a single, engineering-oriented decision criterion [7].

Table 1 contrasts the most relevant directions in the literature with the approach taken in this study, highlighting the specific features enforced in our evaluation pipeline.

Table 1. Related work vs. our contributions (conservative marking).

Ref.	Focus	Multi-HRS Coord.	Distributed Learning (FL/Gossip)	Comm. Realism (Dropout/Bytes)	Robust Aggregation (Byzantine)	Safety Rule (Pressure $\leq \tau$)
[4]	On-site HRS operational planning	—	—	—	—	—
[5]	HRS safety risk assessment	—	—	—	—	risk metrics (not selection rule)
[6,7]	FL systems and edge constraints (surveys)	—	✓	⊖	⊖	—
[8]	Real-time HRS control (learning-based)	—	—	—	—	—
[10,11]	Robust FL (aggregation/Byzantine)	—	✓	—	✓	—
[12,13,15]	Communication reduction (compression/event-triggered)	—	✓	✓/⊖	—	—
[14]	Distributed HRS scheduling (MARL)	✓	—	—	—	—
Ours	Safety-first, communication-realistic, robust distributed scheduling	✓	✓	✓	✓ (where applicable)	✓

Notes: ✓ supported; ⊖ partially/implicitly considered; — not addressed.

Accordingly, this study makes the following contributions:

- **Safety-first feasibility screening for method selection.** We formalize an explicit accept/reject rule based on a pressure feasibility threshold τ and evaluate methods using feasibility-first screening prior to performance ranking.
- **Communication-realistic comparison of centralized and distributed paradigms.** We compare centralized, local-only, and distributed coordination approaches—including server-based FL (FedAvg, FedProx) and peer-to-peer gossip—under heterogeneous NSRDB-driven renewable variability and multi-station demand with realistic participation and delay assumptions.
- **Byte-level communication accounting alongside operational outcomes.** We report operational metrics together with an explicit communication footprint (transmitted bytes) for coordination paradigms with message exchange, enabling transparent safety–performance–communication trade-off assessment.

1.3. Paper Layout

The remainder of this paper is organized as follows: Section 2 describes the system setting and operational variables used in the simulator and defines the performance and safety metrics. Section 3 presents the coordination and learning configurations, including centralized, local-only, gossip, and FL baselines, as well as robustness and communication options. Section 4 reports the results and ablation analyses under the feasibility-first decision criterion. Finally, Section 5 concludes the paper and outlines the practical implications and directions for future work.

2. Model Structure

We consider a multi-station hydrogen service network where stations operate locally yet coordinate without sharing raw operational traces. This motivates federated learning (FL) under data-locality and edge constraints [16]. In the server-based mode, stations act as clients that periodically transmit model updates to a central aggregator; the server aggregates them into a global model and broadcasts the updated parameters to participating stations [17]. Because stations face heterogeneous renewable availability and demand, we explicitly consider the non-i.i.d. regime commonly treated in federated optimization for heterogeneous networks [18].

As a serverless alternative, we also consider a gossip-style coordination overlay where stations exchange updates directly with neighbors over a communication graph, following classical randomized gossip behavior [19]. This decentralized mode inherits topology- and mixing-dependent behavior. Performance varies with the number of gossip exchanges per coordination round under communication constraints [20]. To avoid purely synthetic assumptions in exogenous inputs, renewable signals are taken from real solar and wind time series provided by the National Solar Radiation Database (NSRDB) and mapped to Iraqi locations, naturally inducing heterogeneous operating conditions across stations over time [21].

Figure 1 summarizes the station-level abstraction used in this study, and Figure 2 depicts the two coordination topologies (server-based FL and gossip P2P) as commonly implemented and discussed in the federated learning literature [22].

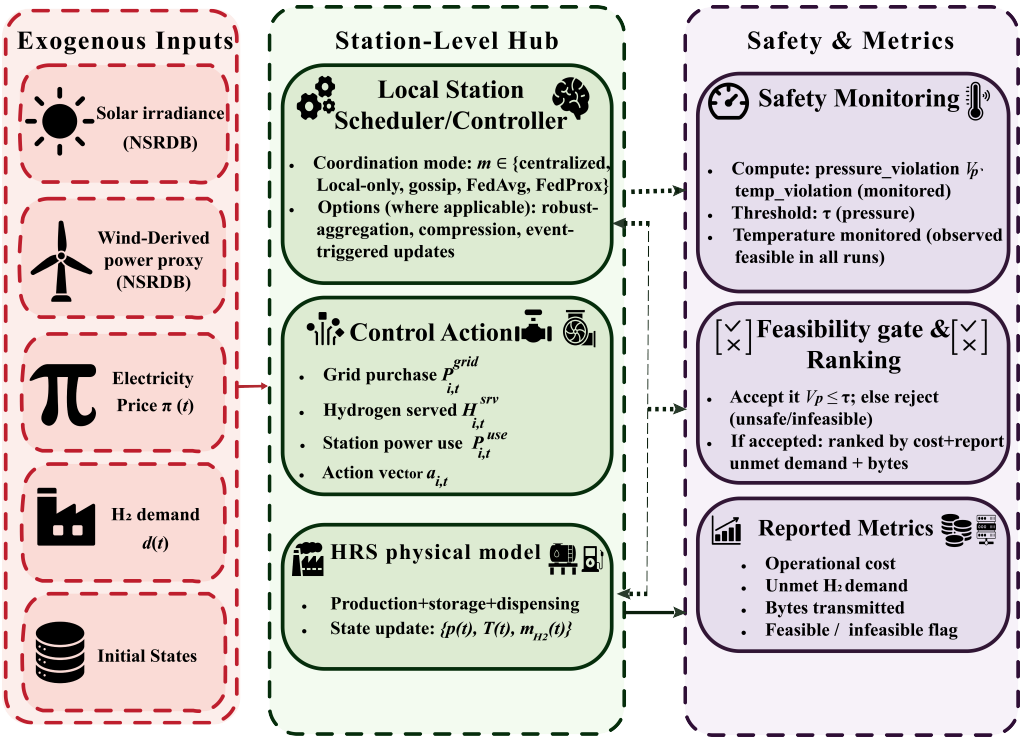


Figure 1. Station-level HRS scheduling pipeline under renewable variability. Symbols used in the diagram are defined as follows: ✓ indicates an accepted or feasible outcome, × indicates rejection due to safety violation, and arrows represent information or control flow between system components. At each time step, the station observes exogenous inputs (e.g., solar/wind availability, demand, and price signals) and local state variables. The controller π_θ selects an operational action (e.g., grid–renewable allocation and service scheduling). The simulator updates the station dynamics and reports performance metrics (operational cost and unmet hydrogen demand) and safety indicators (pressure and temperature violation proxies). A feasibility gate enforces the pressure threshold τ, and only methods satisfying the safety rule are eligible for performance ranking.

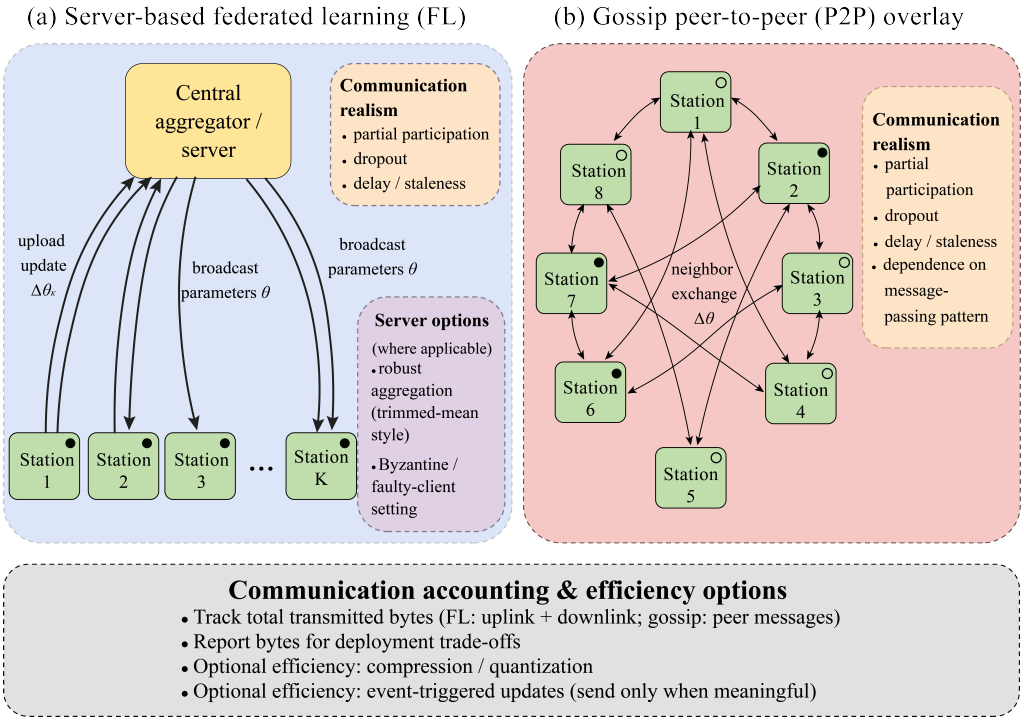


Figure 2. Coordination architectures and communication options used in this study. **(a) Server-based FL:** stations upload local updates (e.g., $\Delta\theta$) to a central aggregator which performs aggregation and

broadcasts model parameters θ each round. Optional server-side robustness includes trimmed-mean aggregation and Byzantine/faulty-client settings. **(b) Gossip P2P:** stations exchange updates with neighbors over an overlay graph (neighbor mixing), making performance dependent on the message-passing pattern. In both paradigms, we model communication realism (partial participation, dropout, and delay/staleness) and account for communication cost in transmitted bytes (FL uplink + downlink; gossip peer messages) with optional efficiency mechanisms such as compression/quantization and event-triggered updates. Symbols used in the diagram are defined as follows: hollow circles represent participating stations, while solid circles represent aggregation or coordination nodes.

Station-Level System Description

Each hydrogen refueling station (HRS) receives electrical power from three sources: solar, wind, and the grid, denoted by $P_{\text{sun}}(t)$, $P_{\text{wind}}(t)$, and $P_{\text{grid}}(t)$, respectively. These exogenous inputs, together with local demand and station state variables, are processed by a local controller (policy π_θ) that computes an operational decision at each time step (e.g., grid–renewable allocation and station operating schedule for serving hydrogen demand) [14].

Following the feasibility-first decision rule introduced in Section 3, a method is considered acceptable only if its aggregate pressure feasibility metric satisfies the safety threshold (pressure violation $\leq \tau$). Performance comparisons are then conducted only within the feasible set. Among feasible methods, we rank by total operating cost and service quality (unmet hydrogen demand). In all reported experiments, pressure feasibility was the binding safety limiter, whereas temperature remained feasible across the tested scenarios (temperature violation $V_T = 0$).

The simulator produces the aggregate metrics required by the reporting pipeline—`cost`, `unmet`, `pressure violation` (V_p), and `temp violation` (V_T)—which are later used to select the best feasible method under a prescribed threshold τ.

Limitations of the pressure proxy. The monitored pressure $p_i(t)$ used in this work is a scheduling-level proxy produced by the station simulator and represents a single aggregated pressure indicator compared against an allowable limit. This abstraction enables controlled comparison across coordination paradigms under identical operating scenarios; however, it does not constitute a full equipment-level digital twin. In real hydrogen stations, feasibility also depends on coupled phenomena not represented here, including spatial thermal gradients and compressor constraints such as ramp-rate limits, startup/shutdown transients, and minimum on/off operation.

Grid interface stress considerations. While feasibility screening in this study is driven by hydrogen-side pressure compliance, the grid exchange $P_{\text{grid}}(t)$ can exhibit fluctuations and ramps under renewable intermittency and control actions. In practical deployments, such variability can contribute to electrical and thermal stress in grid interface equipment (e.g., transformers and regulating devices) through heating and accelerated aging mechanisms. Transformer insulation and thermal considerations are discussed in the transformer material and oil modification review literature [23]. Accordingly, the reported feasibility outcomes should be interpreted as feasibility under the modeled hydrogen-side constraints and monitored proxy, not as a comprehensive certification of grid interface asset health.

Network-level coordination and communication model. Figure 2 illustrates the learning and coordination layers. In the server-based FL mode, participating stations exchange updates with a central aggregator. In the gossip mode, stations exchange updates directly with neighbors, avoiding a single point of failure but introducing dependence on message-passing patterns and mixing behavior.

To reflect deployment realism, the coordination loop can be executed under partial participation and dropout, and communication is tracked as total transmitted bytes. When

Byzantine or faulty-client behavior is enabled, aggregation is evaluated under robustness-oriented settings consistent with Byzantine-robust FL formulations [10].

When communication reduction is enabled, we also evaluate mechanisms that reduce bandwidth while preserving learning utility, such as compression/distillation-style approaches [13].

Event-triggered communication is supported as an additional mechanism to reduce unnecessary transmissions by sending updates only when a trigger condition is satisfied [15].

The above model structure defines the station abstraction, safety monitoring, and coordination layers used throughout the experiments. The next section formalizes the decision rule and distributed learning/scheduling variants evaluated under the feasibility-first protocol.

3. Operating Framework for Safety-Feasible Distributed Scheduling

We consider a network of N hydrogen production stations operating under heterogeneous renewable availability and local demand. Each station executes a local control policy to determine operating setpoints using its own measurements and exogenous inputs (renewables, prices, and demand). In this paper, *scheduling* refers to the time-sequenced selection of station operating setpoints, while the learned policy π_θ acts as a feedback controller that generates these setpoints online from local state measurements. In cyber–physical operations, constraint violations are unacceptable; therefore, safety feasibility must dominate any performance comparison. This safety-first framing is consistent with the broader literature on safety-critical AI systems [24] and learning-enabled safe control paradigms [25]. In addition, verification and assurance frameworks for learning-enabled cyber–physical systems motivate separating safety screening from performance ranking when guarantees are required [26], and control barrier function methods provide a concrete safety-critical control mechanism aligned with admissibility-based decision rules [27].

We evaluate multiple coordination mechanisms, including server-based federated learning (FL) and peer-to-peer gossip schemes. FL is widely used as a privacy-preserving alternative to centralized training under heterogeneous, non-i.i.d. local data [28]. Communication-efficient FL variants are designed to reduce bandwidth while retaining learning stability [29]. Robust FL methods further target failure/adversarial sensitivity in edge environments [30]. Decentralized (gossip-style) approaches remove the central server entirely, trading simplicity for topology- and mixing-dependent behavior [31]. The overall evaluation pipeline is summarized in Figure 3.

To contextualize the mechanisms evaluated in this section, we additionally draw on recent results in decentralized and robust federated learning: multiple gossip steps under communication constraints [32], trimmed-mean aggregation for robustness [33], communication-efficient Byzantine-robust learning with error feedback [34], event-triggered federated communication to reduce unnecessary transmissions [35], and latency-aware semi-synchronous client selection/aggregation in wireless FL [36]. We also cite an MPC-oriented FL survey for general background [37] and a recent validation loss-based decentralized model-sharing approach as an example of server-free coordination [38]. For the hydrogen station safety context, quantitative risk modeling studies motivate strict screening against unsafe operating regimes [39]. The canonical server-based FL baseline used here is Federated Averaging (FedAvg) [40].

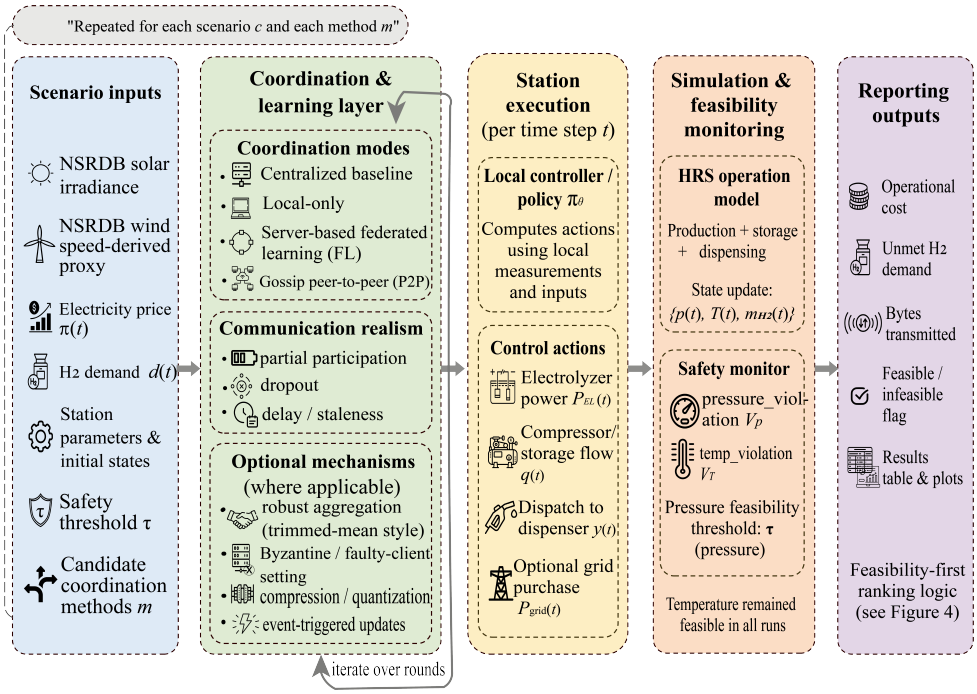


Figure 3. End-to-end evaluation pipeline used in this study. Exogenous scenario traces (renewable generation, demand, and electricity prices) drive the station simulator, while each coordination mode produces a deployed policy π_θ. From the resulting trajectories, we compute operating cost J, unmet demand U, safety metrics (pressure violation V_p and temperature violation V_T), and the communication footprint (total transmitted bytes). Methods are first screened by the feasibility gate ($V_p \leq \tau$) and then ranked by cost within the feasible set.

The feasibility-first ranking logic is summarized in Figure 4.

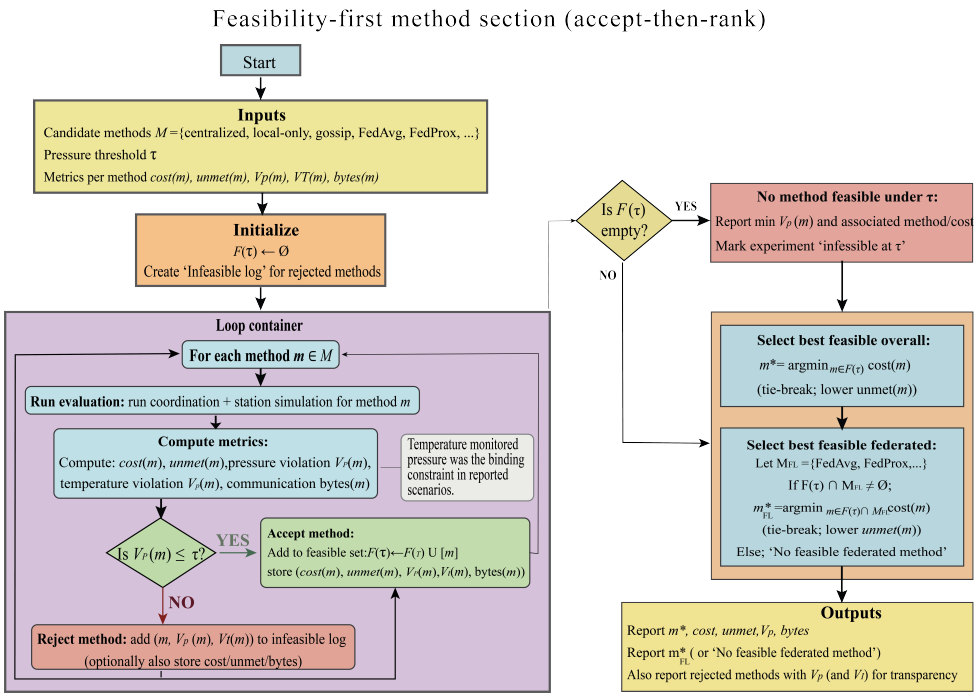


Figure 4. Feasibility-first accept-then-rank decision logic. For each candidate method m, the simulator yields an aggregate pressure violation score $V_p(m)$ and an operating cost $J(m)$. Methods with $V_p(m) > \tau$ are rejected, forming the feasible set $\mathcal{F}(\tau) = \{m \mid V_p(m) \leq \tau\}$. The selected method is $m^\star = \arg\min_{m \in \mathcal{F}(\tau)} J(m)$. If $\mathcal{F}(\tau) = \emptyset$, the case is reported as infeasible (optionally alongside $\min_m V_p(m)$ for transparency).

3.1. Objective Function and Safety Feasibility Decision Rule

Let $i \in \{1, \ldots, N\}$ index stations and $t \in \{1, \ldots, T\}$ index discrete time steps. Each station observes a local state $s_{i,t}$ and applies an action $a_{i,t} = \pi_\theta(s_{i,t})$, where π_θ is a parameterized policy (shared or not, depending on the coordination mode).

The operational objective tracks two practical outcomes: (i) total operating cost and (ii) service quality via unmet hydrogen demand [2]. We define the per-station per-step loss as in Equation (1):

$$\ell_{i,t} = c_t^{\text{grid}} \, [P_{i,t}^{\text{grid}}]_+ + \lambda_u U_{i,t} \tag{1}$$

where $P_{i,t}^{\text{grid}}$ is the net grid exchange (import when positive), $[x]_+ = \max(0, x)$, c_t^{grid} is the time-varying electricity price, $U_{i,t} \geq 0$ is unmet hydrogen demand, and λ_u is the penalty weight [3]. The total objective accumulates the per-step loss as in Equation (2):

$$J(\theta) = \sum_{i=1}^{N} \sum_{t=1}^{T} \ell_{i,t}. \tag{2}$$

Unmet demand is computed from demand $D_{i,t}$ and served hydrogen $H_{i,t}^{\text{srv}}$ as in Equation (3):

$$U_{i,t} = \max\{0, \; D_{i,t} - H_{i,t}^{\text{srv}}\}. \tag{3}$$

Safety is evaluated via pressure feasibility, which is the binding safety limiter in our experiments (temperature remained feasible across the tested scenarios) [5]. Let $p_{i,t}$ denote the monitored pressure proxy and $p^{\max}$ its allowable limit. We measure aggregate pressure violation as in Equation (4):

$$V_p = \frac{1}{NT} \sum_{i=1}^{N} \sum_{t=1}^{T} \max\left\{0, \; \frac{p_{i,t} - p^{\max}}{p^{\max}}\right\}. \tag{4}$$

Safety is enforced via a feasibility-first screening step: Equation (5) filters out any method with $V_p(m) > \tau$ (unsafe), and Equation (6) then performs cost-optimal selection over the remaining feasible candidates using $J(m)$.

$$\mathcal{F}(\tau) = \{ \, m \mid V_p(m) \leq \tau \, \}, \tag{5}$$

where m denotes a method (centralized, local-only, gossip, FedAvg, FedProx, etc.). The selected method is

$$m^\star = \arg \min_{m \in \mathcal{F}(\tau)} J(m). \tag{6}$$

If $\mathcal{F}(\tau) = \emptyset$, no method is deemed safe under the chosen threshold, and results should be reported as infeasible together with $\min_m V_p(m)$ for transparency.

3.2. Station-Level System Model and Operational Constraints

Each station evolves according to a discrete-time simulator as in Equation (7):

$$s_{i,t+1} = f_i(s_{i,t}, a_{i,t}, \xi_{i,t}), \tag{7}$$

where $\xi_{i,t}$ collects exogenous signals such as renewable availability and demand. The simulator outputs operational quantities including grid exchange $P_{i,t}^{\text{grid}}$, hydrogen served $H_{i,t}^{\text{srv}}$, and monitored pressure $p_{i,t}$, as defined in Equation (8):

$$\left(P_{i,t}^{\text{grid}}, \; H_{i,t}^{\text{srv}}, \; p_{i,t}\right) = g_i(s_{i,t}, a_{i,t}, \xi_{i,t}). \tag{8}$$

Safety is enforced by the hard pressure constraint in Equation (9), consistent with station safety limits [5],

$$p_{i,t} \leq p^{\max}, \quad \forall i, t, \tag{9}$$

and violations are summarized by V_p in the previous subsection. Electrical coupling is abstracted through the power balance relation in Equation (10), allowing curtailment/export implicitly through $P_{i,t}^{\text{grid}}$ [4]:

$$P_{i,t}^{\text{grid}} + P_{i,t}^{\text{sun}} + P_{i,t}^{\text{wind}} = P_{i,t}^{\text{use}}, \tag{10}$$

Hydrogen service respects demand satisfaction as stated in Equation (11):

$$0 \leq H_{i,t}^{\text{srv}} \leq D_{i,t}. \tag{11}$$

Justification (feasibility-first vs. soft penalties). The accept-then-rank rule is designed for safety-critical operation where pressure compliance is an acceptability condition rather than a quantity to trade against cost. Importantly, the rule is not "zero-violation"; the threshold τ explicitly allows small aggregate violations (i.e., $V_p \leq \tau$) while rejecting unsafe methods. Soft penalty approaches (e.g., Lagrangian relaxation) require selecting penalty multipliers whose values are scenario-dependent and can change the ranking across methods; moreover, such penalties do not guarantee feasibility when evaluated under non-stationary renewables and communication realism (partial participation, dropout, staleness). Therefore, we use a method-agnostic feasibility gate for safety screening and only then rank cost within the feasible set.

3.3. Distributed Coordination Modes

We evaluate all coordination families under the same station simulator and the same renewable/demand traces, isolating the effect of coordination and communication conditions.

3.3.1. Centralized Baseline

Centralized training pools all station trajectories and optimizes a single policy using global information. This provides an upper-bound reference for learning efficiency, but it is often unrealistic in practice due to privacy constraints and bandwidth bottlenecks [28].

3.3.2. Local Control (Local-Only/No-Communication Baseline)

Each station i learns its own scheduling policy parameters θ_i from local trajectories only, with no communication. The local update rule is given in Equation (12):

$$\theta_i \leftarrow \theta_i - \eta \nabla F_i(\theta_i), \tag{12}$$

where $F_i(\cdot)$ is the station's empirical objective induced by the per-step loss in Equation (1) and the accumulated objective in Equation (2), and η is the local learning rate.

3.3.3. Server-Based Federated Learning (FedAvg/FedProx)

At global round k, a subset of stations S_k participates. Each selected station performs local updates initialized at the broadcast model θ^k:

$$\theta_i^{k,0} = \theta^k, \qquad \theta_i^{k,e+1} = \theta_i^{k,e} - \eta \nabla F_i\Big(\theta_i^{k,e}\Big), \tag{13}$$

which follows the standard FL local training loop in Equation (13) [40]. FedAvg aggregates the locally updated parameters as

$$\theta^{k+1} = \sum_{i \in S_k} w_i \, \theta_i^{k,E}, \qquad w_i = \frac{n_i}{\sum_{j \in S_k} n_j}, \tag{14}$$

where n_i is the number of local samples/steps used at station i (Equation (14)) [40].

FedProx adds a proximal regularization term to mitigate client drift under heterogeneity:

$$F_i^{\text{prox}}(\theta) = F_i(\theta) + \frac{\mu}{2} \left\| \theta - \theta^k \right\|_2^2, \tag{15}$$

which introduces the proximal regularization in Equation (15) to mitigate client drift under heterogeneity [19].

3.3.4. Peer-to-Peer Gossip Coordination

In gossip coordination, stations exchange parameters with neighbors $\mathcal{N}(i)$ using a mixing matrix W,

$$\theta_i^{k+1} = \sum_{j \in \mathcal{N}(i) \cup \{i\}} W_{ij} \, \theta_j^k, \tag{16}$$

removing the server but introducing topology-dependent convergence and sensitivity to packet loss and staleness (Equation (16)) [32].

3.4. Communication Realism and Robustness Options

To reflect practical deployments, we incorporate communication constraints and robustness settings that directly affect learning stability and operational outcomes [28].

Partial participation and dropouts

Client participation is round-dependent; we denote the participating set at round k by S_k, as formalized in Equation (17) [36]:

$$S_k \subseteq \{1, \ldots, N\}. \tag{17}$$

A participating station may fail to deliver an update due to dropout, delay, or staleness; such events are explicitly recorded in the experiment logs and reflected in the effective aggregation set.

Byte-level communication accounting

The total communication footprint is tracked as the cumulative uplink and downlink payload across rounds, as defined in Equation (18):

$$B_{\text{total}} = \sum_k \left(B_k^{\downarrow} + B_k^{\uparrow} \right), \tag{18}$$

where $B_k^{\downarrow}$ and $B_k^{\uparrow}$ denote the downlink and uplink transmitted bytes at round k, respectively.

Compression (optional)

When enabled, each local model update $\Delta\theta_i^k$ is compressed by an operator $Q(\cdot)$ before transmission:

$$\widetilde{\Delta\theta}_i^k = Q\left(\Delta\theta_i^k \right). \tag{19}$$

Quantization/compression is a standard mechanism for reducing uplink load in bandwidth-limited FL deployments (Equation (19)) [34].

Event-triggered communication (optional)

A station transmits an update only if its magnitude exceeds a preset threshold ε:

$$\text{send}_i^k = \mathbb{I}\left\{ \left\| \Delta\theta_i^k \right\|_2 \geq \varepsilon \right\}, \tag{20}$$

where $\mathbb{I}\{\cdot\}$ is the indicator function [32].

Byzantine/faulty clients and robust aggregation (optional)

To tolerate corrupted or faulty updates, the server aggregation can be replaced with a coordinate-wise trimmed mean, as defined in Equation (21) [33]. For coordinate d, let $\{u_{i,d}\}_{i \in S_k}$ be the received values. After sorting and trimming a fraction α, we compute

$$\text{TrimMean}_\alpha(\{u_{i,d}\}) = \frac{1}{|S_k| - 2b} \sum_{r=b+1}^{|S_k|-b} u_{(r),d}, \qquad b = \lfloor \alpha |S_k| \rfloor, \tag{21}$$

where $u_{(r),d}$ is the r-th order statistic.

3.5. Evaluation Metrics and Reporting

For each method m, we report (i) total operating cost $J(m)$ from Equation (2), (ii) total unmet hydrogen demand from Equation (3), (iii) pressure violation $V_p(m)$ from Equation (4), and (iv) total communication bytes $B_{\text{total}}(m)$ from Equation (18). The final method selection follows the feasibility-first decision rule in Equations (5) and (6).

4. Simulation Results and Discussion

The proposed feasibility-first distributed scheduling framework is evaluated in a multi-station hydrogen production setting in which each station operates under renewable variability and station-level safety constraints. All simulations were executed using our Python-based simulator on an Apple MacBook (M3 CPU, 8 GB unified memory, SSD), with the learning components implemented in PyTorch 2.2.1 (MPS acceleration when available). Because the study focuses on comparative algorithmic behavior rather than wall-clock performance, the underlying hardware does not affect the reported comparisons.

We compared (i) centralized coordination, (ii) Local Control (local-only, no coordination) (iii) peer-to-peer gossip coordination, and (iv) server-based federated scheduling (FedAvg and FedProx). For+ each method, we report total cost, aggregate unmet hydrogen demand, safety feasibility indicators, and communication footprint (total transmitted bytes), computed by logging every payload exchanged during coordination.

4.1. Data Description

Renewable inputs are constructed from real-world solar and wind resource time series obtained from the National Solar Radiation Database (NSRDB), which provides hourly irradiance and meteorological fields suitable for operational scheduling studies [21]. In this study, NSRDB data from 2017–2019 are used. In the simulator, solar availability is derived from the NSRDB Global Horizontal Irradiance (GHI) field. When wind is enabled, the simulator also reads the NSRDB wind speed field and converts it into a normalized wind power proxy using a cubic mapping with clipping/saturation (i.e., wind speed→synthetic wind power availability). This proxy preserves the nonlinear wind–power relationship while avoiding unrealistic extremes in the synthesized signal. For consistency with Section 3, the resulting station-indexed renewable inputs are denoted by $r_{i,t}^{\text{sol}}$ and $r_{i,t}^{\text{win}}$, corresponding to the renewable components $P_{\text{sun}}(i,t)$ and $P_{\text{wind}}(i,t)$ used in the power balance model in Equation (10).

To represent heterogeneous operating conditions across geographically separated stations, we introduce a tunable solar–wind mixing parameter $w \in [0,1]$ and form the combined renewable availability signal in Equation (22):

$$r_{i,t} = w\, r_{i,t}^{\mathrm{sol}} + (1-w)\, r_{i,t}^{\mathrm{win}}. \tag{22}$$

Larger values of w yield solar-dominant regimes, whereas smaller values produce wind-dominant regimes. This construction enables controlled stress tests under different renewable dominance regimes while preserving the realistic temporal structure of the NSRDB series. Figure 5 illustrates representative mixed-renewable availability profiles for two regimes ($w = 0.80$ and $w = 0.20$).

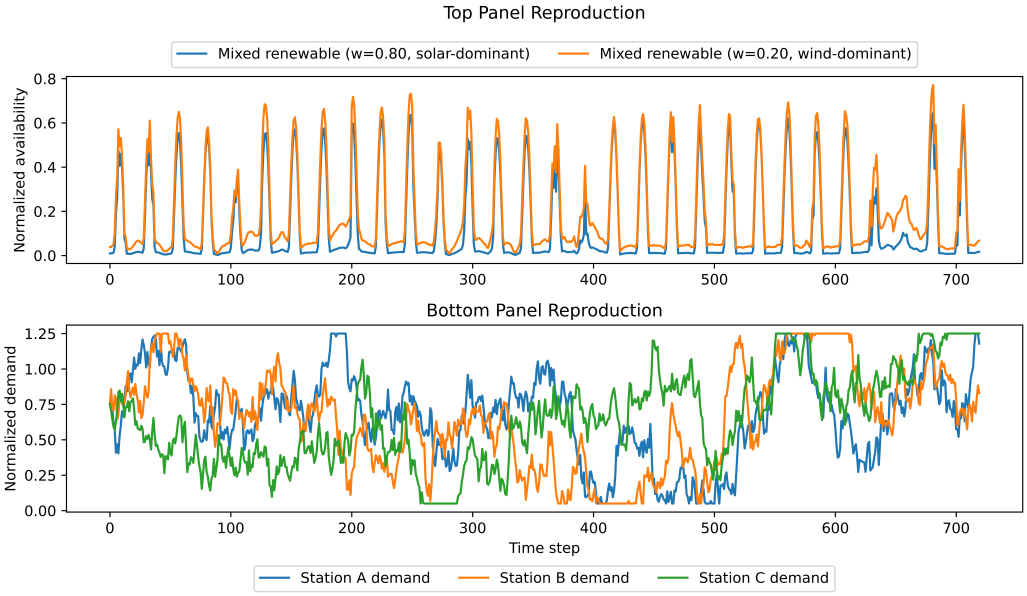


Figure 5. Representative simulator inputs. (**Top**): mixed-renewable availability $r_{i,t}$ formed from NSRDB-derived solar and wind signals for two regimes ($w = 0.80$ solar-dominant vs. $w = 0.20$ wind-dominant). (**Bottom**): heterogeneous station demand traces from the controlled synthetic generator. Within each experiment, these exogenous traces are held fixed across all compared methods to isolate the effect of coordination.

Hydrogen demand and electricity price traces are generated synthetically using a controlled stochastic signal generator and are treated as exogenous inputs to the station simulator. For each station i, the generator is initialized with a deterministic station-specific pseudo-random seed $\texttt{seed}_i = \texttt{seed} + 997 \cdot \texttt{station_id}_i$, which ensures reproducibility while yielding heterogeneous realizations across stations. The generator returns time series $\{r_{i,t}\}_{t=1}^{T}$, $\{D_{i,t}\}_{t=1}^{T}$, and $\{c_{i,t}\}_{t=1}^{T}$ over the simulation horizon T, with variability controlled by noise-scale parameters $(\sigma_r, \sigma_D, \sigma_c)$. When real renewable traces are available, $r_{i,t}$ is overridden by the external dataset (wrap-around repetition if shorter than T), while $D_{i,t}$ and $c_{i,t}$ remain fixed synthetic sequences. Across all coordination methods, the same generated exogenous traces are reused within each experiment so that performance differences reflect coordination and control effects rather than changes in the input realizations.

For transparency, the synthetic demand process is summarized by descriptive statistics computed on the exact traces used in the experiments: mean and standard deviation across stations, coefficient of variation, peak-to-average ratio, and short-lag autocorrelation. In addition, the demand–renewable association is reported using the correlation and cross-correlation between $D_{i,t}$ and the mixed-renewable signal $r_{i,t}$, including the dominant lag if non-zero. These summary statistics are reported in Appendix A (Table A2). Unless

replaced by real measurements, demand and price traces are held fixed across all compared methods to ensure fair comparison and to isolate the effect of coordination and learning from exogenous demand and market variability.

4.2. Case Study

To mirror a practical validation style, we report results from two case studies. Unless stated otherwise, experiments used $N = 10$ stations/clients, a simulation horizon $H = 720$ steps with simulator time step $dt = 1$ (in simulator units), and 100 coordination rounds for methods requiring iterative coordination. To incorporate communication realism, coordinated methods use partial participation (participation fraction per round = 0.60), an independent client dropout probability of 0.15, and a maximum update delay of three rounds.

FedProx tuning protocol (diagnostic μ-sweep): FedProx introduces the proximal coefficient μ in Equation (15). To assess whether the observed infeasibility is configuration-dependent (rather than a reporting artifact), we performed a focused diagnostic μ-sweep under the same communication realism assumptions used throughout the study (partial participation, dropout, bounded delay) and the same scenario inputs (NSRDB traces, mixing weight w, and fixed exogenous demand/price seeds). To reduce computational cost, this diagnostic sweep was executed for 30 coordination rounds (the main comparisons use 100 rounds). For each tested μ, we report cost, unmet demand, and pressure violation, and we mark feasibility using the same gate as the main paper (pressure violation $\leq \tau$). The diagnostic outcomes are summarized in Appendix A (Table A1) to make the μ-sensitivity transparent and reproducible. The FedProx entry reported in Table 2 corresponds to the 100-round main comparison configuration.

Case 1 (Feasibility-first baseline comparison). Centralized, Local Control, gossip, and federated methods are compared under the same renewable profiles and station constraints, with feasibility enforced through a τ-based acceptability rule ($\tau = 0.2$).

Case 2 (Renewable-dominance stress test via mixing weights). The solar–wind mixing parameter w is varied to evaluate the effect of renewable dominance on feasibility and cost. In particular, we contrast low-cost but unsafe configurations (infeasible due to pressure violation) with the best feasible configuration under $\tau = 0.2$.

Table 2. Feasibility-first comparison for $\tau = 0.2$.

Method/Tag	Coordination	Cost	Unmet	Temp Violation V_T	Pressure Violation V_p	Total Bytes	Feasible ($\tau = 0.2$)
centralized_real	Centralized	3396.25	100.39	0.00	0.867	N/A	×
localonly_real	Local Control	2131.83	59.32	0.00	0.172	0	✓
gossip_real	Gossip (P2P)	2339.52	66.19	0.00	0.284	710,500	×
fedavg_real_solarwind_w0.20	FedAvg (FL)	2423.72	68.71	0.00	0.163	866,688	✓
fedavg_real_solarwind_w0.15	FedAvg (FL)	1952.27	46.74	0.00	2.631	1,079,808	×
fedprox_real_solarwind_w0.20	FedProx (FL)	2033.53	51.68	0.00	1.505	1,761,792	×

Notes: ✓ indicates feasible ($V_p \leq \tau$); × indicates infeasible ($V_p > \tau$); N/A indicates not applicable (no message-based coordination for the centralized baseline).

4.3. Scheduling Results Analysis

4.3.1. Cost and Service Performance Under Feasibility Filtering

Table 2 summarizes the key outcomes of the feasibility-first rule ($\tau = 0.2$). The centralized baseline yields a high operational cost and violates pressure feasibility (pressure violation $V_p = 0.867$), indicating that it is rejected under the acceptability rule despite representing a traditional centralized coordination baseline.

Among feasible methods, the Local Control baseline achieves the lowest cost (2131.83) with pressure violation $V_p = 0.172$, corresponding to a 37.23% reduction relative to the

centralized baseline (3396.25). The Local Control baseline also reduces unmet demand from 100.39 to 59.32 (a 40.91% reduction), indicating that feasibility-first selection does not necessarily sacrifice service quality.

For federated coordination, the best feasible configuration under $\tau = 0.2$ is FedAvg with solar–wind mixing $w = 0.20$, achieving a cost of 2423.72, unmet demand of 68.71, and pressure violation $V_p = 0.163$. Peer-to-peer gossip coordination remains near-feasible but fails the $\tau = 0.2$ gate (pressure violation $V_p = 0.284$), and FedProx is strongly infeasible in this scenario (pressure violation $V_p = 1.505$).

Figures 6 and 7 show the evolution of cost and unmet demand over coordination rounds for the iterative methods.

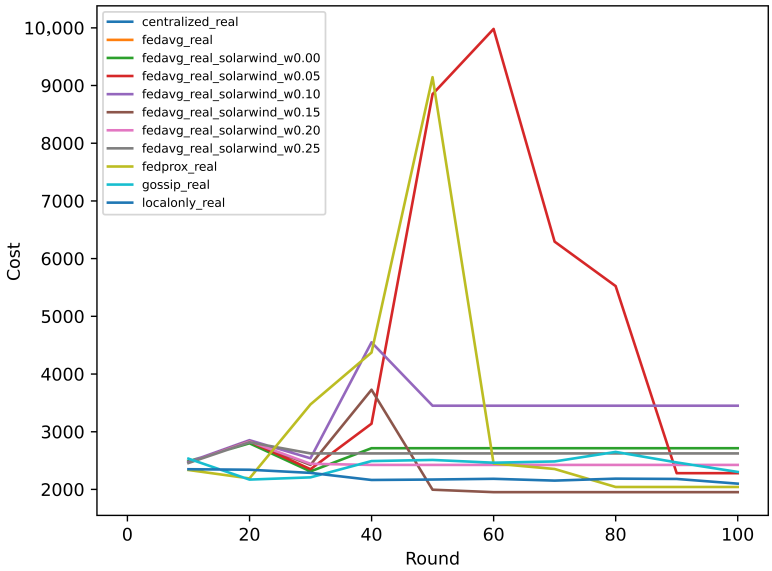


Figure 6. Total operating cost versus coordination round for iterative coordination methods (federated and gossip). At each round, the currently deployed policy is evaluated under the same simulator inputs (identical scenario traces and station constraints), and the resulting cost J is logged. The curves therefore summarize convergence behavior and final cost performance under the chosen communication realism settings.

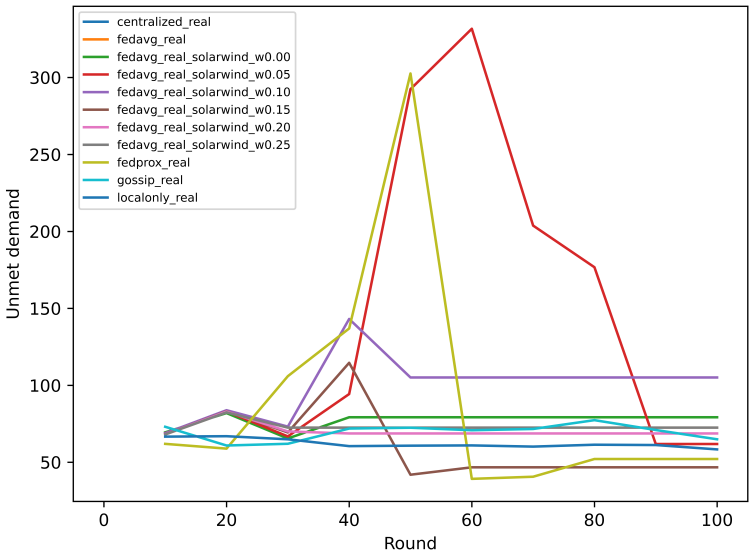


Figure 7. Aggregate unmet hydrogen demand versus coordination round for iterative coordination methods (federated and gossip). Values are logged under the same simulator inputs used for the cost traces, showing how coordination affects service quality over rounds in addition to cost.

4.3.2. Pressure Feasibility Behavior (Safety Dominates Selection)

A key observation is that feasibility, not raw cost, determines acceptability. The lowest-cost configuration in this study is not necessarily usable. For example, FedAvg with $w = 0.15$ achieved a cost of 1952.27 (and unmet demand of 46.74), but it was infeasible due to a large pressure violation (pressure violation $V_p = 2.631$) and was therefore rejected. This directly validates the feasibility-first ranking rule: a cost-only selection would choose an unsafe policy. Soft penalty (Lagrangian-style) selection would face the same failure mode unless penalty multipliers are tuned to enforce the same safety threshold, which is scenario-dependent; therefore, we retain an explicit feasibility gate for auditable acceptability (Section 3). Figure 8 shows the pressure violation metric V_p over coordination rounds and visualizes the τ cut-off.

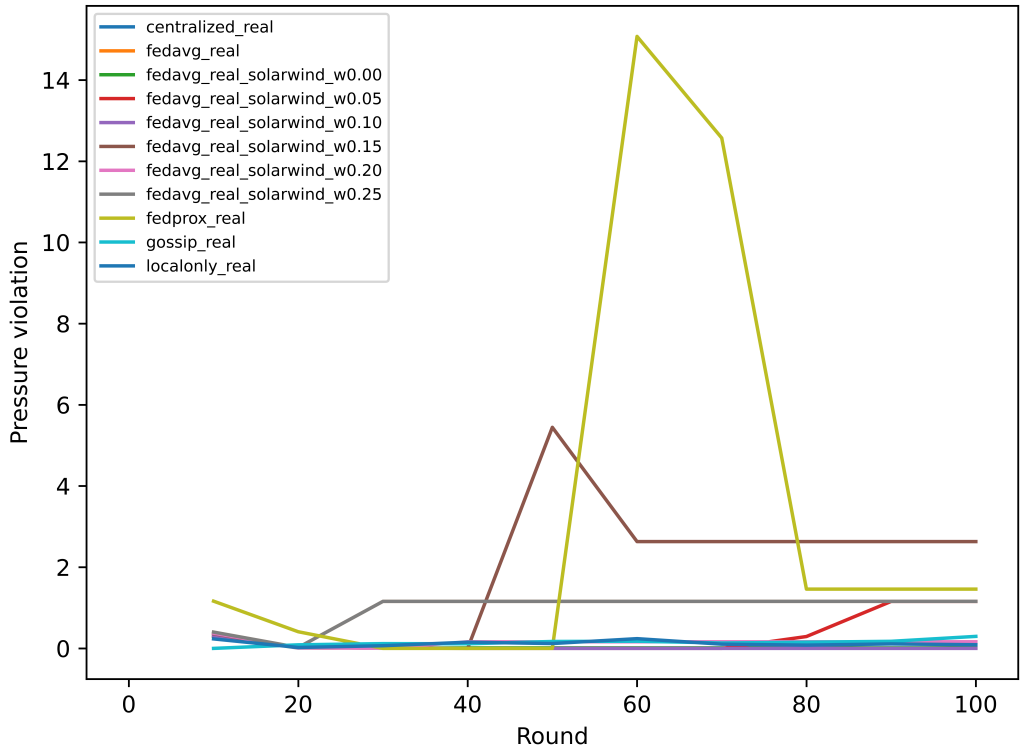


Figure 8. Pressure violation metric V_p versus coordination round for the iterative coordination methods. Curves show V_p for each method across rounds. The feasibility threshold used in the study is $\tau = 0.2$.

4.3.3. Communication Footprint and Practicality

Communication overhead is a tangible deployment cost and must be reported explicitly alongside operational outcomes. In our experiments, the best feasible federated configuration (FedAvg, $w = 0.20$) incurred a total communication footprint of 866,688 bytes, whereas the Local Control baseline required no communication by design. Gossip coordination exchanged 710,500 bytes in total, yet it still failed the feasibility gate ($\tau = 0.2$), indicating that increased communication alone does not guarantee safety feasibility.

For the centralized baseline, communication is not directly comparable because centralized optimization implicitly assumes global access to all station trajectories, signals, and constraints. To avoid an apples-to-oranges comparison, Table 2 reports Total Bytes only for coordination methods in which message payloads are explicitly exchanged and logged. Figure 9 visualizes the cumulative communication load (bytes transmitted) across coordination rounds, highlighting the non-trivial overhead required by the iterative methods relative to Local Control operation.

The reported Total Bytes are the aggregate serialized coordination payload exchanged over all coordination rounds (Appendix A.2). For the 100-round experiments, the corresponding average payload is on the order of 7.1 kB/round (gossip), 8.7 kB/round (FedAvg, $w = 0.20$), 10.8 kB/round (FedAvg, $w = 0.15$), and 17.6 kB/round (FedProx). In practical terms, such payloads are lightweight for broadband backhaul links, but they can be constraining on low-rate industrial IoT links. For example, LPWAN technologies such as LoRaWAN operate at low data rates with additional duty-cycle and payload-size constraints; therefore, exchanging multi-kilobyte payloads repeatedly can translate to seconds of airtime per round (and minutes over a full 100-round run), making frequent coordination impractical unless coordination rounds are infrequent and/or strong compression/sparsification is used. By contrast, narrowband cellular IoT (e.g., NB-IoT/LTE-M) typically supports substantially higher effective uplink throughput than LPWAN, so the same per-round payloads can be transmitted with much smaller time overhead, making the communication–optimality trade-off more realistic for coordinated deployments. Overall, these results motivate reporting both feasibility and communication cost: coordination can improve performance, but its practicality depends on the available link budget and the desired coordination frequency.

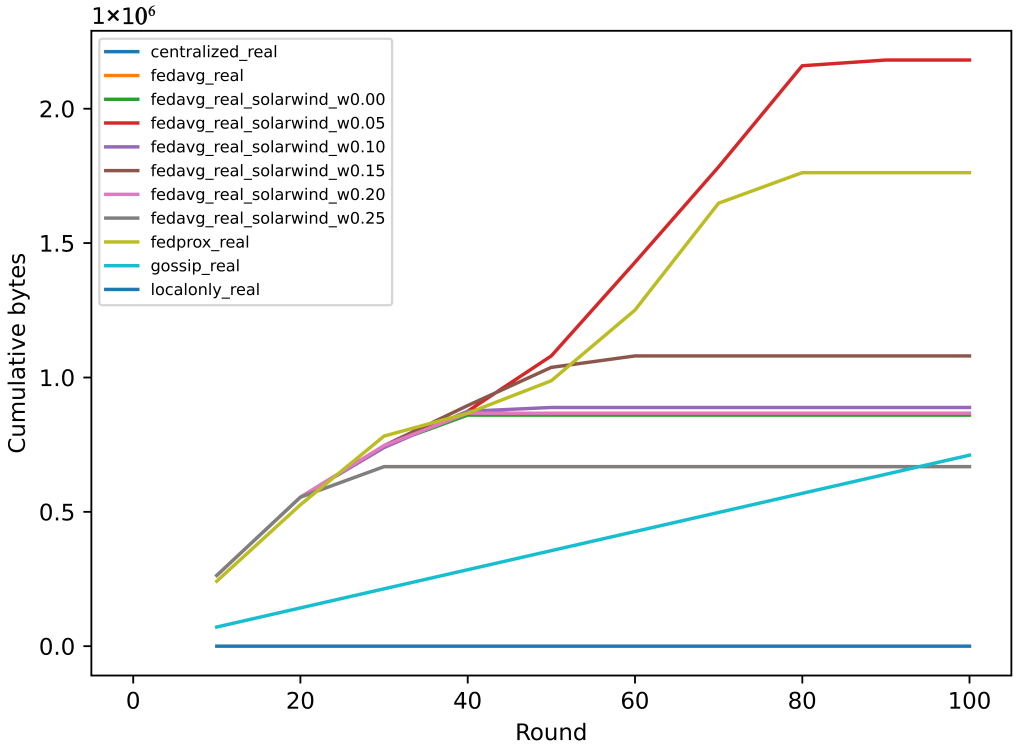


Figure 9. Cumulative transmitted bytes versus coordination round for iterative coordination methods. Bytes are accumulated from the serialized uplink+downlink payloads (server-based FL) or peer-to-peer payloads (gossip) recorded during execution, providing a deployment-facing communication cost alongside performance and feasibility.

4.4. Performance–Safety–Communication Benefit Analysis

4.4.1. Cost Benefit

Under $\tau = 0.2$, Local Control is the lowest-cost feasible strategy and outperforms centralized coordination by 37.23%. Among the coordinated methods, the best feasible federated configuration (FedAvg, $w = 0.20$) remains competitive ($J = 2423.72$) while achieving improved pressure feasibility behavior relative to the centralized baseline.

4.4.2. Safety Benefit

The feasibility-first gate prevents the selection of unsafe low-cost policies (e.g., FedAvg $w = 0.15$) which would otherwise appear superior under cost-only ranking. This provides a deployment-style selection criterion in which acceptability is not negotiable. While soft penalties can be used, they require choosing penalty weights that may vary across operating regimes; our explicit τ-gated screening avoids this tuning dependence and ensures that only policies meeting the safety criterion are compared by cost.

4.4.3. Communication Practicality

By explicitly reporting transmitted bytes, coordination claims become measurable rather than qualitative. The results demonstrate a clear performance–safety–communication trade-off; coordination can improve service behavior and feasibility outcomes, but it introduces communication dependence that must be budgeted and engineered in practical deployments.

4.5. Validation Analysis

4.5.1. Method Comparison Under τ-Based Acceptability

Table 2 provides the defensible comparison under $\tau = 0.2$. The centralized baseline is rejected due to feasibility failure; Local Control is the best feasible overall; and FedAvg ($w = 0.20$) is the best feasible federated configuration. Infeasible runs are still reported for transparency and to demonstrate the necessity of feasibility-first screening.

4.5.2. Consistency with Hydrogen Station Optimization and FL System Practice

Prior hydrogen refueling station scheduling studies commonly emphasize cost and service quality, but they often do not enforce an explicit feasibility-first accept/reject criterion when comparing competing coordination paradigms. Similarly, many federated learning studies report learning or operational performance without byte-level communication accounting. In contrast, this study formalizes a deployment-style decision rule (screen for safety feasibility, then rank by performance) and treats transmitted bytes as a first-class metric. Table 3 contrasts the proposed framework against representative related work in terms of feasibility screening, communication accounting, and safety modeling.

Table 3. Feature-level comparison with representative related work.

Work	Hard Feasibility Screening	Reports Bytes	Multi-Station Distributed Coordination	Renewable Variability	Explicit Pressure/Temperature Safety
Representative HRS scheduling (centralized) [2]	Partial/implicit	No	No	Often yes	Often yes
Distributed energy scheduling without hard screening [4]	No	No	Sometimes	Yes	Sometimes
Federated learning systems evaluation practice [19]	N/A	Yes	Yes	N/A	N/A
This work	Yes (τ-gated)	Yes (byte-level logging)	Yes (FedAvg/FedProx/Gossip)	Yes (NSRDB-driven)	Yes (pressure and temperature)

Notes: N/A indicates not applicable to the referenced work or not reported/defined in that study.

5. Conclusions

This study evaluated distributed coordination for multi-station hydrogen production under NSRDB-driven renewable variability using an explicit safety-first decision rule. Instead of ranking methods by cost alone, we applied a feasibility gate ($V_p \leq \tau$) and

compared methods only within the feasible set while also reporting unmet demand and communication overhead (Total Bytes) to make deployment trade-offs transparent.

The results demonstrate that feasibility can overturn cost-only selection: the lowest-cost configuration (FedAvg, $w = 0.15$; cost 1952.27) is not deployable because it violates the pressure threshold ($V_p = 2.631 > \tau = 0.2$). Within the feasible set, the best federated configuration (FedAvg, $w = 0.20$) satisfies the safety gate ($V_p = 0.163$) with cost 2423.72 and unmet demand 68.71, requiring 866,688 transmitted bytes. Local Control remains a strong baseline and achieves the best feasible overall outcome (cost 2131.83, unmet 59.32, $V_p = 0.172$), indicating that coordination should be adopted only when it yields clear improvements under the same feasibility screening.

Future work will increase realism by incorporating measured demand and dynamic prices, expanding stress tests across participation and delay regimes, and integrating communication-reduction mechanisms (e.g., compression and event-triggered updates) while preserving auditable feasibility-first screening.

Author Contributions: M.T.: conceptualization, methodology, software, investigation, data curation, visualization, formal analysis, writing—original draft. S.W.A.-S.: supervision, methodology, validation, formal analysis, writing—review and editing. All authors have read and agreed to the published version of the manuscript.

Funding: This research received no external funding.

Data Availability Statement: The data and code used to support the findings of this study are available from the corresponding author upon reasonable request.

Conflicts of Interest: The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Nomenclature

The following abbreviations are used in this manuscript:

A. Acronyms

HRS	Hydrogen refueling station (hydrogen production/serving station in the simulations)
NSRDB	National Solar Radiation Database (source of renewable traces)
GHI	Global Horizontal Irradiance (solar irradiance field used from NSRDB)
FL	Federated learning
FedAvg	Federated Averaging coordination method
FedProx	Federated Proximal coordination method
P2P	Peer-to-peer (gossip) communication
IID	Independent and identically distributed (data)
OTA	Over-the-air update transmission

B. Indices and Sets

i	Station index, $i = 1, \ldots, N$
t	Time step index, $t = 1, \ldots, T$
k	Communication round index (federated/gossip coordination)
e	Local update/epoch index within a round
m	Method/policy index (e.g., centralized, local-only, FedAvg, FedProx, gossip)
S_k	Set of participating stations in round k
$\mathcal{N}(i)$	Neighbor set of station i in the gossip topology
d	Model parameter coordinate index (used in robust aggregation)
N	Number of stations
T	Scheduling horizon length (number of time steps)

C. Signals, State, and Operational Quantities

$s_{i,t}$	Local station state at station i and time t
$a_{i,t}$	Scheduling/control action applied at station i and time t
$\xi_{i,t}$	Exogenous inputs at station i and time t (renewable availability and demand signals)
$f_i(\cdot)$	Station state transition function
$g_i(\cdot)$	Station output map producing operational outputs from $(s_{i,t}, a_{i,t}, \xi_{i,t})$
$P_{i,t}^{\text{sun}}$	Solar-derived electrical power available at station i, time t (NSRDB-driven)

$P_{i,t}^{\text{wind}}$	Wind speed-derived proxy power available at station i, time t (NSRDB-driven)
$P_{i,t}^{\text{grid}}$	Net grid power exchange at station i, time t (import when positive)
$P_{i,t}^{\text{use}}$	Power consumption of the electrolyzer/production process at station i, time t
$H_{i,t}$	Hydrogen storage level at station i, time t
$H_{i,t}^{\text{srv}}$	Hydrogen served/produced (delivered to demand) at station i, time t
$D_{i,t}$	Hydrogen demand at station i, time t
$p_{i,t}$	Monitored station pressure proxy at station i, time t
$p^{\max}$	Maximum allowable pressure proxy (station safety limit)
w	Solar–wind mixing weight used for renewable-dominance stress tests
D. Objective, Metrics, and Feasibility Rule	
$\ell_{i,t}$	Per-step operational cost at station i, time t
c_t^{grid}	Time-varying electricity price at time t
$U_{i,t}$	Unmet hydrogen demand at station i, time t
λ_u	Penalty weight applied to unmet demand in the objective
$J(\theta)$	Total objective for policy parameters θ
$\text{cost}(m)$	Total operating cost achieved by method m
$\text{unmet}(m)$	Aggregate unmet hydrogen demand achieved by method m
V_p	Aggregate pressure safety violation metric
$V_T(m)$	Aggregate temperature violation metric for method m
τ	Feasibility threshold for pressure violations ($V_p \leq \tau$)
$\mathcal{F}(\tau)$	Feasible set of methods under threshold τ (accept–then–rank stage)
B_{total}	Total communication volume (bytes) exchanged across all coordination rounds
$B_k^{\uparrow}$	Uplink bytes transmitted in round k
$B_k^{\downarrow}$	Downlink bytes transmitted in round k
E. Learning/Coordination and Communication Accounting	
$\pi_\theta(\cdot)$	Parameterized local policy/controller mapping states to actions
θ	Model/policy parameter vector
θ_i^k	Local parameters at station i during/after round k
θ^k	Reference/global parameters at round k
η	Learning rate
$F_i(\theta)$	Local objective function at station i
μ	FedProx proximal coefficient
E	Number of local epochs/steps per coordination round
n_i	Local sample/trajectory count used for weighting

Appendix A. Reproducibility Addendum: Feasibility Gate, Communication Accounting, and FedProx Sensitivity

This appendix consolidates implementation details added during revision to improve reproducibility and to clarify (i) the feasibility-first evaluation rule, (ii) byte-level communication accounting, and (iii) the FedProx hyperparameter sensitivity behavior.

Appendix A.1. Feasibility-First Gate

All compared methods and configurations are evaluated using the same acceptability rule (consistent with Equation (5)). A method/configuration is considered feasible if and only if

$$V_p(m) \leq \tau, \tag{A1}$$

where $V_p(m)$ is computed using Equation (4) from the pressure trajectory $\{p_{i,t}\}$ generated when running method m, and $\tau = 0.2$ for the main feasibility-first comparison. The feasibility gate is applied before cost ranking; only feasible candidates are eligible for cost-based selection (Equation (6)).

Appendix A.2. Communication Accounting (Total Bytes)

For coordinated methods (server-based federated learning and peer-to-peer gossip), *Total Bytes* is reported as the cumulative serialized coordination payload size summed over all coordination rounds (consistent with Equation (18)). This accounts for explicit coordination messages required by each paradigm.

- **Local-only:** Total Bytes = 0 (no coordination messages).
- **Centralized baseline:** Total Bytes = N/A, because it assumes a single globally informed controller rather than message-based coordination among distributed stations.
- **Federated learning/gossip (including FedProx):** Total Bytes is computed by summing the per-round transmitted payload sizes recorded in the run outputs.

Appendix A.3. FedProx Hyperparameter Sensitivity (Diagnostic)

FedProx modifies the local objective by adding a proximal regularizer (Equation (15)) controlled by $\mu \geq 0$, which constrains local updates toward the broadcast model. To assess whether the observed infeasibility is configuration-dependent, we ran a focused diagnostic μ-sweep under the *same scenario inputs and communication realism settings* as the main FedProx experiment (same NSRDB traces and mixing weight w, participation fraction, dropout probability, staleness bound, and the same random seeds/exogenous traces). To reduce computational cost, each diagnostic run is executed for 30 coordination rounds (the main comparisons use 100 rounds); therefore, the sweep is used to document sensitivity rather than to determine the main table FedProx configuration.

Evaluation logic (consistent with feasibility-first reporting)

For each μ, we report cost, unmet, and the aggregate pressure violation metric V_p, and we mark feasibility using the same criterion as the main paper ($V_p \leq \tau$, with $\tau = 0.2$). If no tested μ meets the feasibility threshold, the sweep still identifies which values minimize V_p under the fixed experimental setting. Note that $\mu = 0$ is included as a sanity check (the proximal term vanishes and the objective reduces to the FedAvg form under the same participation schedule).

Table A1. FedProx μ-sweep (diagnostic sensitivity) under the same scenario inputs/settings as the main FedProx experiment; each run executed for 30 coordination rounds.

μ	Cost	Unmet	Temp Violation	Pressure Violation V_p	Feasible ($\tau = 0.2$)
0.00	1996.45	54.52	0.00	0.0168	Yes
0.01	1895.55	49.99	0.00	0.5232	No
0.10	2150.91	58.86	0.00	0.2105	No
1.00	4609.64	145.26	0.00	0.0000	Yes

Interpretation: Feasibility under the strict τ-gate is μ-sensitive; some μ values increase pressure violation beyond τ, whereas stronger proximal regularization can suppress violations but may degrade operational performance (cost/unmet). This diagnostic supports the main discussion that algorithmic stabilization alone does not guarantee feasibility when constraints are enforced through an external feasibility-first screening rule.

Appendix A.4. Synthetic Demand Characterization

Hydrogen demand traces are synthetically generated and held fixed across all compared methods within each experiment to ensure that observed differences arise from coordination and control rather than from changes in exogenous inputs. Table A2 summarizes descriptive statistics computed on the exact demand traces used in the simulator, together with empirical association metrics relative to the mixed-renewable availability signal.

Table A2. Statistical characterization of synthetic hydrogen demand (mean ± std across stations).

μ_D	σ_D	CV $= \sigma_D/\mu_D$	PAR	ρ_1	corr(D, r)	Dominant Lag (Steps)
0.582 ± 0.086	0.364 ± 0.060	0.645 ± 0.161	2.186 ± 0.292	0.978 ± 0.010	-0.192 ± 0.162	8.3 ± 39.8

Columns: mean demand (μ_D), standard deviation (σ_D), coefficient of variation (CV), peak-to-average ratio (PAR), lag-1 autocorrelation (ρ_1), Pearson correlation (corr(D, r)), and dominant cross-correlation lag between demand and renewable availability (if non-zero).

Appendix A.5. Reproducibility Settings Summary

All fixed simulation, communication, and training settings used in the experiments are consolidated in Table A3.

Table A3. Reproducibility settings used across experiments.

Setting	Value
Number of stations/clients N	10
Simulation horizon T (time steps)	720
Time step dt	1 (simulator units)
Coordination rounds K (main comparisons)	100
Participation fraction per round	0.60
Dropout probability per round	0.15
Maximum update delay (staleness bound)	3 rounds
Feasibility threshold τ	0.2
Solar–wind mixing weight w (main case)	0.20
NSRDB years used	2017–2019
Station seed rule	$\texttt{seed}_i = \texttt{seed} + 997 \cdot \texttt{station_id}_i$
Local epochs/steps E	2
Learning rate η	0.001
Batch size	128
FedProx proximal coefficient (main run)	0.02
FedProx μ-sweep rounds (diagnostic)	30

References

1. Jiang, L.; Fu, X. An Ammonia–Hydrogen Energy Roadmap for Carbon Neutrality: Opportunity and Challenges in China. *Engineering* **2021**, *7*, 1688–1691. [CrossRef]
2. Cardona, P.; Costa-Castelló, R.; Roda, V.; Carroquino, J.; Valiño, L.; Ocampo-Martinez, C.; Serra, M. Modelling and Operation Strategy Approaches for On-Site Hydrogen Refuelling Stations. *Int. J. Hydrogen Energy* **2024**, *52*, 49–64. [CrossRef]
3. Parsibenehkohal, R.; Jamil, M.; Khan, A.A. A Multi-Stage Framework for Coordinated Scheduling of Networked Microgrids in Active Distribution Systems with Hydrogen Refueling and Charging Stations. *Int. J. Hydrogen Energy* **2024**, *71*, 1442–1455. [CrossRef]
4. Fang, X.; Wang, Y.; Dong, W.; Yang, Q.; Sun, S. Optimal Energy Management of Multiple Electricity–Hydrogen Integrated Charging Stations. *Energy* **2023**, *262*, 125624. [CrossRef]
5. Kwak, H.; Kim, M.; Min, M.; Park, B.; Jung, S. Assessing the Quantitative Risk of Urban Hydrogen Refueling Station in Seoul, South Korea, Using SAFETI Model. *Energies* **2024**, *17*, 867. [CrossRef]
6. Fekri, M.N.; Patel, H.; Grolinger, K.; Sharma, V. Deep Learning for Load Forecasting with Smart Meter Data: Online Adaptive Recurrent Neural Network. *Appl. Energy* **2021**, *282*, 116177. [CrossRef]
7. Lim, W.Y.B.; Luong, N.C.; Hoang, D.T.; Jiao, Y.; Liang, Y.-C.; Yang, Q.; Niyato, D.; Miao, C. Federated Learning in Mobile Edge Networks: A Comprehensive Survey. *IEEE Commun. Surv. Tutor.* **2020**, *22*, 2031–2063. [CrossRef]
8. Huy, T.H.B.; Duy, N.T.M.; Phu, P.V.; Le, T.-D.; Park, S.; Kim, D. Robust Real-Time Energy Management for a Hydrogen Refueling Station Using Generative Adversarial Imitation Learning. *Appl. Energy* **2024**, *373*, 123847. [CrossRef]
9. Hegedűs, I.; Danner, G.; Jelasity, M. Decentralized Learning Works: An Empirical Comparison of Gossip Learning and Federated Learning. *J. Parallel Distrib. Comput.* **2021**, *148*, 109–124. [CrossRef]
10. Ma, X.; Zhou, Y.; Wang, L.; Miao, M. Privacy-Preserving Byzantine-Robust Federated Learning. *Comput. Stand. Interfaces* **2022**, *80*, 103561. [CrossRef]
11. Peng, J.; Wu, Z.; Ling, Q.; Chen, T. Byzantine-Robust Variance-Reduced Federated Learning over Distributed Non-i.i.d. Data. *Inf. Sci.* **2022**, *616*, 367–391. [CrossRef]
12. Pillutla, K.; Kakade, S.M.; Harchaoui, Z. Robust Aggregation for Federated Learning. *IEEE Trans. Signal Process.* **2022**, *70*, 1142–1154. [CrossRef]

13. Wu, C.; Wu, F.; Lyu, L.; Huang, Y.; Xie, X. Communication-Efficient Federated Learning via Knowledge Distillation. *Nat. Commun.* **2022**, *13*, 2032. [CrossRef] [PubMed]
14. Jiang, Y.; Liu, J.; Zheng, H. Optimal Scheduling of Distributed Hydrogen Refueling Stations for Fuel Supply and Reserve Demand Service with Evolutionary Transfer Multi-Agent Reinforcement Learning. *Int. J. Hydrogen Energy* **2024**, *54*, 239–255. [CrossRef]
15. Xue, X.; Mao, H.; Li, Q.; Huang, F.; Abd El-Latif, A.A. An Energy Efficient Specializing DAG Federated Learning Based on Event-Triggered Communication. *Mathematics* **2022**, *10*, 4388. [CrossRef]
16. Shlezinger, N.; Chen, M.; Eldar, Y.C.; Poor, H.V.; Cui, S. UVeQFed: Universal Vector Quantization for Federated Learning. *IEEE Trans. Signal Process.* **2021**, *69*, 500–514. [CrossRef]
17. Yang, Q.; Liu, Y.; Chen, T.; Tong, Y. Federated Machine Learning: Concept and Applications. *ACM Trans. Intell. Syst. Technol.* **2019**, *10*, 1–19. [CrossRef]
18. Xia, Q.; Ye, W.; Tao, Z.; Wu, J.; Li, Q. A Survey of Federated Learning for Edge Computing: Research Problems and Solutions. *High-Confid. Comput.* **2021**, *1*, 100008. [CrossRef]
19. Kairouz, P.; McMahan, H.B. Advances and Open Problems in Federated Learning. *Found. Trends Mach. Learn.* **2021**, *14*, 1–210. [CrossRef]
20. Liu, B.; Lv, N.; Guo, Y.; Li, Y. Recent Advances on Federated Learning: A Systematic Survey. *Neurocomputing* **2024**, *597*, 128019. [CrossRef]
21. Sengupta, M.; Xie, Y.; Lopez, A.; Habte, A.; Maclaurin, G.; Shelby, J. The National Solar Radiation Database (NSRDB). *Renew. Sustain. Energy Rev.* **2018**, *89*, 51–60. [CrossRef]
22. Li, T.; Sahu, A.K.; Zaheer, M.; Sanjabi, M.; Talwalkar, A.; Smith, V. Federated Optimization in Heterogeneous Networks. *Proc. Mach. Learn. Syst. (MLSys)* **2020**, *2*, 429–450.
23. Boyd, S.; Ghosh, A.; Prabhakar, B.; Shah, D. Randomized Gossip Algorithms. *IEEE Trans. Inf. Theory* **2006**, *52*, 2508–2530. [CrossRef]
24. Perez-Cerrolaza, J.; Abella, J.; Borg, M.; Donzella, C.; Cerquides, J.; Cazorla, F.J.; Englund, C.; Tauber, M.; Nikolakopoulos, G.; Flores, J.L. Artificial Intelligence for Safety-Critical Systems in Industrial and Transportation Domains: A Survey. *ACM Comput. Surv.* **2024**, *56*, 1–40 . [CrossRef]
25. Dawson, C.; Gao, S.; Fan, C. Safe Control With Learned Certificates: A Survey of Neural Lyapunov, Barrier, and Contraction Methods for Robotics and Control. *IEEE Trans. Robot.* **2023**, *39*, 1749–1767. [CrossRef]
26. Tran, H.-D.; Xiang, W.; Johnson, T.T. Verification Approaches for Learning-Enabled Autonomous Cyber–Physical Systems. *IEEE Des. Test* **2022**, *39*, 24–34. [CrossRef]
27. Cohen, M.H.; Molnar, T.G.; Ames, A.D. Safety-Critical Control for Autonomous Systems: Control Barrier Functions via Reduced-Order Models. *Annu. Rev. Control* **2024**, *57*, 100947. [CrossRef]
28. Li, T.; Talwalkar, A.; Smith, V. Federated Learning: Challenges, Methods, and Future Directions. *IEEE Signal Process. Mag.* **2020**, *37*, 50–60. [CrossRef]
29. Zong, R.; Qin, Y.; Wu, F.; Tang, Z.; Li, K. FedCS: Efficient Communication Scheduling in Decentralized Federated Learning. *Inf. Fusion* **2024**, *102*, 102028. [CrossRef]
30. Zhou, H.; Zheng, Y.; Jia, X. Towards Robust and Privacy-Preserving Federated Learning in Edge Computing. *Comput. Netw.* **2024**, *243*, 110321. [CrossRef]
31. Tang, Z.; Shi, S.; Li, B.; Chu, X. GossipFL: A Decentralized Federated Learning Framework With Sparsified and Adaptive Communication. *IEEE Trans. Parallel Distrib. Syst.* **2023**, *34*, 909–922. [CrossRef]
32. Hashemi, A.; Acharya, A.; Das, R.; Vikalo, H.; Sanghavi, S.; Dhillon, I. On the Benefits of Multiple Gossip Steps in Communication-Constrained Decentralized Federated Learning. *IEEE Trans. Parallel Distrib. Syst.* **2022**, *33*, 2727–2739. [CrossRef]
33. Wang, T.; Zheng, Z.; Lin, F. Federated Learning Framework Based on Trimmed Mean Aggregation Rules. *Expert Syst. Appl.* **2025**, *270*, 126354. [CrossRef]
34. Ghosh, A.; Maity, R.K.; Kadhe, S.; Mazumdar, A.; Ramchandran, K. Communication-Efficient and Byzantine-Robust Distributed Learning With Error Feedback. *IEEE J. Sel. Areas Inf. Theory* **2021**, *2*, 942–953. [CrossRef]
35. Lu, S.; Gao, Z.; Zhang, P.; Xu, Q.; Xie, T.; Zhang, A. Event-Triggered Federated Learning for Fault Diagnosis of Offshore Wind Turbines With Decentralized Data. *IEEE Trans. Autom. Sci. Eng.* **2024**, *21*, 1271–1283. [CrossRef]
36. Yu, L.; Sun, X.; Albelaihi, R.; Yi, C. Latency-Aware Semi-Synchronous Client Selection and Model Aggregation for Wireless Federated Learning. *Future Internet* **2023**, *15*, 352. [CrossRef]
37. Liu, F.; Zheng, Z.; Shi, Y.; Tong, Y.; Zhang, Y. A Survey on Federated Learning: A Perspective from Multi-Party Computation. *Front. Comput. Sci.* **2023**, *18*, 181336. [CrossRef]
38. Mamba Kabala, D.; Hafiane, A.; Bobelin, L.; Canals, R. Decentralized Federated Learning Using Validation Loss for Model Sharing in Crop Disease Classification. *Ecol. Inform.* **2025**, *90*, 103205. [CrossRef]

39. Liang, Z.; Yang, Y.; Wang, Y.; Zhang, M.; Zhuang, Y. Multidimensional Quantitative Modeling Fusion Analysis of Safety Risks in Hydrogen Refueling Stations: A Case Study of a Station in Beijing. *J. Loss Prev. Process Ind.* **2026**, *100*, 105865. [CrossRef]
40. McMahan, B.; Moore, E.; Ramage, D.; Hampson, S.; y Arcas, B.A. Communication-Efficient Learning of Deep Networks from Decentralized Data. In *20th International Conference on Artificial Intelligence and Statistics (AISTATS)*; PMLR: Cambridge, MA, USA, 2017; Volume 54, pp. 1273–1282.

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.