

Energy Markets for Machine Intelligence and Cryptographic Settlement

Matthew Long

The YonedaAI Collaboration

YonedaAI Research Collective

Chicago, IL

matthew@yonedaai.com · <https://yonedaai.com>

23 July 2026

Abstract

Large AI facilities and proof-of-work mines are electricity loads, but they are not interchangeable batteries and they are not automatically helpful to a grid. Their economic usefulness depends on location, timing, controllability, telemetry, deadlines, restart losses, network congestion, reserve obligations, and the market rule used to pay them. This paper develops a bounded market model in which heterogeneous compute loads bid for energy on a lossless direct-current network. Generators offer energy and upward reserve. AI and cryptographic workloads declare different minimum load, ramp, deadline, restart, curtailment, and telemetry constraints. The clearing problem maximizes submitted bid surplus subject to nodal balance, line limits, generator capability, reserve procurement, and workload feasibility.

Two accounting results are proved within that model. First, summing nodal balance over a lossless network yields exact system energy balance. Second, settlement at any complete set of nodal prices satisfies an exact identity: load energy payments minus generator energy credits equal line congestion rent. If reserve charges are allocated to equal reserve credits, the combined energy and reserve cash account also closes. These are algebraic results, not claims that a discrete bid clears at a competitive equilibrium price.

A deterministic Python simulator enumerates small discrete schedules and independently checks balance, line limits, reserves, deadlines, and ramps. Tests cover congestion, line and generator outages, telemetry failure, baseline inflation, strategic bids, restart costs, and reserve scarcity. Every fixture is synthetic. The direct-current formulation is useful for active-power dispatch and congestion accounting, but it cannot prove voltage, reactive-power, protection, frequency, or dynamic-stability adequacy. A flexible compute project should be described as grid-responsive only after those omitted engineering questions and the relevant tariff have been checked.

Keywords: flexible load; artificial intelligence; proof of work; locational marginal pricing; demand response; reserve; direct-current optimal power flow; congestion; telemetry; baseline

Evidence status. The balance and settlement identities are PROVED for the stated lossless network and accounting boundary. The dispatch, outage, congestion, telemetry, baseline, and strategic-bid examples are COMPUTATIONAL. Claims about a particular data center, mine, balancing area, tariff outcome, or dynamic-stability benefit are OBSTRUCTED without site-specific data and engineering studies. Empirical estimates of bid curves, response times, restart costs, and counterfactual baselines are OPEN.

Contents

1	Flexible computation is a market participant, not a slogan	5
1.1	Contributions	5
1.2	Nonclaims	6
2	Institutional facts that constrain the model	6
2.1	FERC market description	6
2.2	DOE recommendations for AI and data centers	6
2.3	ERCOT large loads and controllable load resources	7
2.4	CAISO demand response, meters, and telemetry	7
2.5	NERC reliability vocabulary	8
2.6	DC optimal power flow as a limited reference model	8
2.7	Prior computing-flexibility research	8
3	Accounting boundary and market objects	9
3.1	Space, time, and units	9
3.2	Nodes and lines	9
3.3	Generators and reserve	9
3.4	Fixed and flexible demand	10
3.5	Telemetry fallback	10
3.6	Submitted value and audited value	10
4	Heterogeneous compute feasible sets	11
4.1	An AI service is not one resource type	11
4.2	Proof-of-work settlement is also heterogeneous	11
4.3	A comparison contract	11
5	The dispatch and accounting theorems	12
6	Baseline settlement and the counterfactual problem	13
7	The reference implementation	14
7.1	Why exhaustive discrete clearing	14
7.2	Generic data structures	14
7.3	Independent dispatch checks	15
7.4	What the software does not do	15

8	Computational scenarios	15
8.1	Energy and reserve co-constraint	15
8.2	Congestion allocates scarce transfer capability	16
8.3	Nodal congestion-rent identity	16
8.4	Line outage and island balance	17
8.5	Generator outage and reserve infeasibility	17
8.6	AI deadline, ramp, and restart	17
8.7	Telemetry failure	17
8.8	Strategic bid	18
8.9	Baseline inflation	18
9	Adversarial market-design analysis	18
9.1	Baseline manipulation	18
9.2	Telemetry failure and stale control	18
9.3	Strategic bids and market power	19
9.4	Congestion and rebound	19
9.5	Outages and common-mode risk	19
9.6	Reserve double counting	19
10	What a production study must add	20
10.1	Alternating-current feasibility	20
10.2	Dynamic stability	20
10.3	Security-constrained and unit-commitment detail	20
10.4	Distribution and retail terms	20
10.5	Emissions and water	20
11	Market design for responsive compute	21
11.1	Participation product	21
11.2	Direct scheduling versus baseline response	21
11.3	AI queue design	21
11.4	Mining fleet design	21
11.5	Scarcity and emergency operation	22
12	Empirical research program	22
12.1	Minimum dataset	22
12.2	Causal questions	22
12.3	Identification designs	23
12.4	Outcome reporting	23
13	Verification contract	23
13.1	Executable checks	23
13.2	Reproduction	24
13.3	Interpretation of a passing suite	24
14	Limits and failure conditions	24

15 Decision protocol for a real project	25
16 Claim ledger	26
17 Conclusion	26
A Compact mathematical program	27
B Settlement derivation by incidence matrix	27
C Audit schema	27
D Reviewer checklist	28

1 Flexible computation is a market participant, not a slogan

Electricity can be turned into inference, model training, search, database work, cryptographic hashes, and settlement assurance. That common energy input does not make those outputs economically equivalent. It also does not make the loads equally flexible. An inference service may face a latency target that expires in seconds. A training run may have a daily completion deadline but expensive checkpoint and restart behavior. A proof-of-work machine can often stop quickly, yet a site may still have cooling constraints, firmware delays, minimum import obligations, hedges, pool contracts, or restart wear. A useful market representation starts with those differences.

The policy discussion often begins one step too late. It asks whether a data center or mine can “support the grid” before defining the node, interval, service, baseline, control path, and contingency. A load that follows price in an uncongested hour may worsen a local constraint in another hour. A load that offers curtailment from an inflated baseline may receive money without delivering a real reduction. A load whose telemetry fails may be unavailable exactly when an operator needs to observe or control it. A project that passes a lossless active-power calculation may still fail a voltage or stability study.

The right question is narrower:

Under a declared network model and market rule, which compute schedules are physically feasible, which maximize submitted bid surplus, and which settlement identities can be verified from the resulting dispatch?

This paper answers that question with a small model that can be inspected line by line. It does not attempt to reproduce a production system operator. The model has enough structure to expose common category errors:

1. energy balance is not the same as reserve adequacy;
2. reserve headroom is not the same as frequency response;
3. a nodal price is not an average retail tariff;
4. a flexible upper bound is not proof of deliverable curtailment;
5. a reported baseline is not an observed counterfactual;
6. a welfare result based on bids is not a truthful-value result;
7. a direct-current dispatch is not an alternating-current or dynamic stability certificate.

1.1 Contributions

The paper contributes five connected pieces.

First, it gives a common market representation for machine intelligence and proof-of-work settlement without forcing the workloads into one generic “flexible load” shape. Each resource declares its own bounds, ramp limits, deadline energy, restart cost, fallback power under telemetry failure, and submitted value.

Second, it adds those workloads to a node-indexed and interval-indexed lossless dispatch problem with generators, transmission lines, outages, energy bids, reserve offers, and reserve requirements.

Third, it proves active-power balance and nodal settlement identities within the model. These results are simple enough to be checked independently of the optimizer.

Fourth, it implements a deterministic standard-library simulator for small instances. The program chooses a bid-surplus-maximizing schedule by exhaustive enumeration. It records feasible-candidate counts, nodal residuals, line flows, reserve awards, bid surplus, and true surplus.

Fifth, it treats adversarial cases as first-class tests. Baseline gaming, strategic bids, telemetry loss, congestion, line outages, generator outages, reserve scarcity, and impossible deadlines are not footnotes. They are the cases most likely to break an attractive but underspecified proposal.

1.2 Nonclaims

The model does not estimate the elasticity of AI or mining demand. It does not forecast a regional load shape. It does not determine a production locational marginal price. It does not represent unit commitment, minimum generator up and down times, losses, reactive power, voltage, protection, short-circuit strength, harmonics, frequency response, or transient stability. It does not claim that AI is more valuable than proof of work, or the reverse. It does not infer public benefit from private willingness to pay.

2 Institutional facts that constrain the model

2.1 FERC market description

The Federal Energy Regulatory Commission describes organized wholesale markets as systems that use day-ahead and real-time processes to match supply and demand while respecting resource and transmission constraints [1, 2]. Its current market materials distinguish energy from ancillary services and describe economic dispatch, security constraints, and locational marginal prices. In the standard decomposition, an LMP includes energy, congestion, and marginal loss components.

This institutional description creates a boundary for the present paper. A node-specific price must not be replaced by a national average electricity price. Congestion is not an optional narrative adjustment after dispatch. It is a consequence of a binding network constraint. Real-time balancing also does more than settle deviations. It supports the continuous operational task of matching load, generation, and interchange.

The reference simulator retains active-power balance and congestion but removes marginal losses. Its line-rent identity is therefore a lossless identity. Applying it to a market with marginal-loss pricing requires a separate loss account.

2.2 DOE recommendations for AI and data centers

The United States Department of Energy’s 2024 recommendations on powering AI and data-center infrastructure identify uncertainty in the timing, scale, and shape of new loads. The report discusses load flexibility as an opportunity, but it also emphasizes grid cost,

resource adequacy, reliability, planning, and coordination [4]. The report does not establish that every data center is flexible or beneficial.

The useful implication is procedural. A project should declare at least:

- expected and maximum import by interval;
- the fraction that can be curtailed and the time needed to do so;
- workload deadlines and quality consequences;
- restart and recovery behavior;
- on-site generation and storage, if any;
- telemetry and control architecture;
- the relevant interconnection and market participation path.

The model below turns the first five items into constraints. Telemetry appears as an eligibility and fallback condition. Interconnection requirements remain outside the optimizer because they depend on the jurisdiction and facility.

2.3 ERCOT large loads and controllable load resources

ERCOT’s current Large Load Integration materials and Planning Guide include a dedicated process for large-load interconnection and modification [5, 6]. The current materials include pathways and forms for large loads evaluated as provisional controllable load resources. ERCOT’s July 2026 Nodal Protocols define controllable load resources and contain the live rules for dispatch, metering, telemetry, and settlement [7]. The load-resource participation materials require the resource’s qualified scheduling entity to establish specified telemetry points [8].

These are not decorative administrative details. A large load that can change power internally but cannot communicate or satisfy registration rules is not equivalent to a market-qualified resource. Likewise, provisional treatment in an interconnection study is not a finding of dynamic stability. The simulator therefore separates physical fallback power from market controllability.

2.4 CAISO demand response, meters, and telemetry

CAISO allows qualified demand-response resources to participate in day-ahead and real-time energy markets, and eligible resources can provide ancillary services. Its public guidance requires a scheduling coordinator and certified revenue-quality meters. Ancillary-service participation requires additional certification and direct telemetry [9]. CAISO also states that accurate metering supports settlement and that direct telemetry gives the operator real-time visibility into participating generators and loads [10].

CAISO’s Demand Response Business Practice Manual defines demand-response energy measurement using an approved performance-evaluation method. For a customer-load-baseline method, the relevant quantity compares the calculated baseline with actual underlying load during an event [11]. This is why a baseline is an estimand, not a meter reading. The meter observes actual load. It does not observe what the same facility would have consumed without dispatch.

2.5 NERC reliability vocabulary

NERC standards define requirements for planning and operating the North American bulk power system [12]. The current NERC glossary distinguishes frequency error, frequency regulation, frequency response, and frequency-response obligation [13]. These distinctions matter for compute-load claims. A schedule that preserves a five-megawatt reserve margin in a static dispatch has not demonstrated a measured response in megawatts per frequency deviation. Nor has it demonstrated activation time, sustainment, recovery, or deliverability.

In this paper, “reserve” means only upward megawatt headroom procured from declared generator offers for an interval. It is a modeled capacity constraint. It is not a NERC frequency-response finding.

2.6 DC optimal power flow as a limited reference model

Schweppe and coauthors established the location-sensitive spot-pricing frame for electricity systems [14]. Standard power-system texts develop economic dispatch and optimal power flow with network constraints [15]. MATPOWER provides transparent AC and DC power-flow and optimal-power-flow formulations. Its published architecture supports dispatchable loads and branch constraints [16, 17].

The DC approximation is useful because its active-power flow equations are linear. That simplicity also marks its limit. It does not solve alternating current voltage magnitude, reactive power, or nonlinear loss equations. Steady-state DC feasibility is silent about frequency dynamics, controls, protection, fault ride-through, and post-contingency recovery. Every computational result below carries that limitation.

2.7 Prior computing-flexibility research

The idea that computing can respond to electricity conditions predates the current AI-load debate. Liu and coauthors formulated geographical load balancing across data centers and showed that the result depends on the pricing rule and the willingness of services to shift work [18]. Their later demand-response pricing model directly examined data-center market power and prediction error [19]. Wierman and coauthors surveyed the operational and market-design opportunities and challenges for data-center demand response [20]. Yang and Chien’s Zero-Carbon Cloud work examined high-performance computing powered by otherwise uneconomic renewable generation and quantified the cost and workload implications of that volatile supply [21].

Cryptocurrency mining has also been modeled as a location-sensitive, price-responsive load. Menati, Lee, and Xie used a synthetic Texas grid to study mining demand-response participation and found that price effects vary with capacity and location [22]. Menati and coauthors subsequently analyzed reliability, emissions, and market effects with high-resolution synthetic-grid scenarios [23]. Those studies motivate several design choices here, especially nodal congestion, explicit flexibility, and adverse price effects. The present contribution is not the first proposal for flexible computing. Its narrower contribution is a shared, audit-oriented feasible-set model with exact balance and settlement checks and explicit failure cases for telemetry, baselines, reserves, and strategic bids.

3 Accounting boundary and market objects

3.1 Space, time, and units

Let $\mathcal{N}$ be a finite set of nodes and $\mathcal{T}$ a finite ordered set of intervals. Each interval has duration $\Delta > 0$ hours. Power variables use megawatts. Multiplication by Δ produces megawatt-hours. Monetary bids and marginal generation costs use one declared currency unit per megawatt-hour. Reserve offers use the declared currency unit per megawatt-hour of reserved capability, so multiplying an award by interval duration gives the interval reserve cost.

The conversion to joules is exact:

$$1 \text{ MWh} = 3.6 \times 10^9 \text{ J.}$$

The market is written in power-system units because bids and constraints are normally expressed that way. An energy-normalized economic report can divide the resulting value account by joules after the market boundary is closed.

3.2 Nodes and lines

Each line $\ell \in \mathcal{L}$ has an oriented from-node $o(\ell)$, a to-node $d(\ell)$, susceptance $B_\ell > 0$, thermal capacity $\bar{f}_\ell \geq 0$, and outage indicator $a_{\ell t} \in \{0, 1\}$. The orientation is an accounting convention. A negative flow means power moves opposite the declared direction.

The lossless DC flow equation is

$$f_{\ell t} = a_{\ell t} B_\ell (\theta_{o(\ell),t} - \theta_{d(\ell),t}), \quad (1)$$

with

$$-a_{\ell t} \bar{f}_\ell \leq f_{\ell t} \leq a_{\ell t} \bar{f}_\ell. \quad (2)$$

One angle per connected island is fixed to zero. If an outage splits the network, each island must balance independently.

3.3 Generators and reserve

Each generator $g \in \mathcal{G}$ is located at node $n(g)$. It has interval availability a_{gt} , minimum and maximum energy output $\underline{p}_g, \bar{p}_g$, marginal energy cost c_g , maximum upward reserve offer $\bar{r}_g$, and reserve offer cost k_g . The implementation uses

$$a_{gt} \underline{p}_g \leq p_{gt} \leq a_{gt} \bar{p}_g, \quad (3)$$

$$0 \leq r_{gt} \leq \bar{r}_g, \quad (4)$$

$$p_{gt} + r_{gt} \leq a_{gt} \bar{p}_g. \quad (5)$$

For interval reserve requirement R_t ,

$$\sum_{g \in \mathcal{G}} r_{gt} \geq R_t. \quad (6)$$

Reserve is procured from the cheapest eligible offers after each candidate energy schedule is formed. This greedy step is exact for the separable, single-product reserve constraint in eq. (6). It would not be exact for zonal deliverability, multiple response classes, coupled ramping, or nonconvex reserve bids.

3.4 Fixed and flexible demand

Fixed demand at node n and interval t is $d_{nt} \geq 0$. Each flexible workload $w \in \mathcal{W}$ has node $n(w)$ and scheduled power x_{wt} . The nodal balance equation is

$$\sum_{g:n(g)=n} p_{gt} - d_{nt} - \sum_{w:n(w)=n} x_{wt} = \sum_{\ell:o(\ell)=n} f_{\ell t} - \sum_{\ell:d(\ell)=n} f_{\ell t}. \quad (7)$$

Every workload declares minimum and maximum power:

$$\underline{x}_w \leq x_{wt} \leq \bar{x}_w. \quad (8)$$

It also declares initial power $x_{w,-1}$, upward ramp u_w , and downward ramp v_w :

$$-v_w \leq x_{wt} - x_{w,t-1} \leq u_w. \quad (9)$$

When workload w has deadline τ_w and required energy E_w ,

$$\Delta \sum_{t \leq \tau_w} x_{wt} \geq E_w. \quad (10)$$

Let $z_{wt} = 1$ when the schedule moves from zero power to positive power, and let $s_w \geq 0$ be the declared restart cost. In the discrete implementation, the objective subtracts $s_w z_{wt}$. This is a simple restart representation. It does not model a full state machine for checkpointing, thermal recovery, or machine-level boot sequences.

3.5 Telemetry fallback

Let $m_w \in \{0, 1\}$ indicate whether the resource has usable telemetry under the modeled participation rule. When telemetry is required and $m_w = 0$, the market cannot choose any point in eq. (8). The implementation fixes the schedule to a declared physical fallback q_w :

$$x_{wt} = q_w. \quad (11)$$

This choice avoids a dangerous modeling shortcut. Missing telemetry does not make a physical load disappear. It removes the modeled control option.

3.6 Submitted value and audited value

Each workload declares bid value β_w per megawatt-hour and an optional audited or estimated true value v_w . Dispatch uses β_w , not v_w . The discrete clearing objective is

$$\max \left[\Delta \sum_{w,t} \beta_w x_{wt} - \Delta \sum_{g,t} c_g p_{gt} - \Delta \sum_{g,t} k_g r_{gt} - \sum_{w,t} s_w z_{wt} \right]. \quad (12)$$

The ex-post true-surplus audit replaces β_w with v_w while leaving the chosen schedule fixed. A strategic bid can therefore raise submitted surplus while lowering audited true surplus.

4 Heterogeneous compute feasible sets

4.1 An AI service is not one resource type

“AI load” can refer to online inference, batch inference, model training, data preparation, checkpoint writing, retrieval, or a mixed service. An online service may have a tight minimum-power requirement because latency and availability contracts leave little time to defer. A batch queue can have a low interval minimum but a binding cumulative deadline. A training job may be interruptible only at checkpoints and may lose work when restarted.

The model does not encode behavior by the text label “AI.” Instead, a resource is made AI-like by its declared constraints. For example:

- online inference: high $\underline{x}_w$, modest ramp limits, no broad daily deadline substitution;
- batch inference: low $\underline{x}_w$, positive E_w , and a deadline τ_w ;
- training: low minimum power, a deadline, finite ramp rates, and material restart cost s_w ;
- preemptible experimentation: low minimum, loose deadline, small restart cost, and lower bid value.

Quality adjustment remains outside the electricity market unless the bid itself internalizes it. A megawatt-hour that produces low-quality or unused inference should not receive a high economic value merely because the hardware stayed busy.

4.2 Proof-of-work settlement is also heterogeneous

A proof-of-work site’s private energy value depends on block subsidy, fees, hash price, pool terms, hardware efficiency, cooling overhead, curtailment contracts, and financial hedges. Many machines can stop faster than a continuous industrial process. That fact supports a potentially broad curtailment range. It does not imply zero restart cost, perfect telemetry, or unlimited response duration.

A mining resource can be represented with low minimum power, a high downward ramp limit, and a bid derived from expected incremental settlement revenue. A site with contractual minimum import or cooling constraints needs a higher $\underline{x}_w$. A site that takes time to reconnect needs finite upward ramp and positive restart cost. A site seeking reserve or demand-response payments must also satisfy the relevant qualification and measurement rules.

4.3 A comparison contract

Table 1: Feasible-set questions for compute loads

Field	AI examples	Proof-of-work examples
Minimum power	latency floor, availability reserve, memory residency	cooling, network, or contractual import floor
Maximum power	rack, feeder, thermal, accelerator availability	machine fleet, transformers, cooling, interconnection limit

Field	AI examples	Proof-of-work examples
Ramp down	request shedding, queue control, checkpoint interval	firmware control, machine contactors, thermal management
Ramp up	queue recovery, cache warmup, checkpoint restart	staged machine restart, pool reconnection, thermal limit
Deadline	response latency, batch completion, training milestone	usually economic horizon rather than task completion
Restart cost	lost work, recovery delay, reliability risk	wear, downtime, lost pool revenue, recovery overhead
Telemetry fallback	keep critical service on, shed only prequalified queue	continue at site fallback, local controller action, or trip
Bid value	task and quality adjusted service value	expected incremental settlement revenue and risk adjustment

The table is a measurement checklist, not a default parameter set.

5 The dispatch and accounting theorems

Assumption 5.1 (Lossless oriented network). For each interval, line flow is lossless and appears once as an outflow from its from-node and once as an inflow to its to-node. Every node satisfies eq. (7).

Proposition 5.2 (System balance). *Under Assumption 5.1, every feasible interval satisfies*

$$\sum_{g \in \mathcal{G}} p_{gt} = \sum_{n \in \mathcal{N}} d_{nt} + \sum_{w \in \mathcal{W}} x_{wt}. \quad (13)$$

Proof. Sum eq. (7) over all nodes. Each line flow appears once with a positive sign on the right side at its from-node and once with a negative sign at its to-node. The line terms cancel. The remaining left side is total generation minus fixed demand minus flexible demand. Setting it to zero gives eq. (13). $\square$

Remark 5.3. The proposition proves an active-power accounting identity. It does not prove that frequency stays at its schedule after a disturbance. Frequency adequacy requires dynamic response, controls, timing, and contingency analysis that are absent here.

Let λ_{nt} be any declared nodal energy price for every node and interval. Define load energy payment

$$P^D = \Delta \sum_{n,t} \lambda_{nt} \left(d_{nt} + \sum_{w:n(w)=n} x_{wt} \right), \quad (14)$$

generator energy credit

$$C^G = \Delta \sum_{g,t} \lambda_{n(g),t} p_{gt}, \quad (15)$$

and line congestion rent

$$K = \Delta \sum_{\ell, t} (\lambda_{d(\ell), t} - \lambda_{o(\ell), t}) f_{\ell t}. \quad (16)$$

Theorem 5.4 (Lossless nodal settlement identity). *Under Assumption 5.1, every feasible dispatch and every complete set of nodal prices satisfy*

$$P^D - C^G = K. \quad (17)$$

Proof. Multiply each nodal balance equation by $-\Delta\lambda_{nt}$, then sum over nodes and intervals. The generation and demand terms become $P^D - C^G$. For line ℓ , the from-node contribution is $-\Delta\lambda_{o(\ell), t} f_{\ell t}$, while the to-node contribution is $+\Delta\lambda_{d(\ell), t} f_{\ell t}$. Their sum is the corresponding term in K . Summing all lines gives eq. (17). $\square$

Suppose reserve is paid at interval price ρ_t . Total reserve credit is

$$C^R = \Delta \sum_{g, t} \rho_t r_{gt}.$$

If reserve charges P^R are allocated so that $P^R = C^R$, then

$$P^D + P^R = C^G + C^R + K. \quad (18)$$

This extension is an accounting convention. Individual system operators can allocate reserve and uplift through different tariff rules.

Limitation 5.5 (No price-formation theorem). The simulator takes settlement prices as a complete external input. The discrete enumeration in eq. (12) does not establish that those prices are supporting dual variables, that bids are truthful, or that the outcome is a competitive equilibrium.

6 Baseline settlement and the counterfactual problem

Demand response measures a change relative to normal or counterfactual use. FERC's public definition refers to changes from normal consumption patterns in response to prices or incentive payments [3]. That definition contains a counterfactual. Actual interval load is observable with a suitable meter. Normal load in the absence of dispatch is not simultaneously observable.

Let b^{rep} be the reported baseline power, b^{ver} a later verified counterfactual estimate, m metered power, Δ interval duration, and π the payment per megawatt-hour. The claimed reduction is

$$q^{\text{claim}} = \Delta \max\{0, b^{\text{rep}} - m\}, \quad (19)$$

while the verified reduction is

$$q^{\text{verified}} = \Delta \max\{0, b^{\text{ver}} - m\}. \quad (20)$$

The excess payment attributable to baseline difference is

$$X = \pi \max\{0, q^{\text{claim}} - q^{\text{verified}}\}. \quad (21)$$

This calculation does not determine which baseline method is best. It exposes the consequence of a disagreement. If the reported baseline is 10 MW, the verified counterfactual is 6 MW, actual load is 6 MW, the interval is one hour, and the payment is 100 currency units per MWh, the claim is 4 MWh even though the verified reduction is zero. The resulting excess payment is 400 currency units.

Baseline design creates incentives. A participant may have reason to increase load in baseline-setting periods, avoid dispatch that would lower its historical baseline, or select a method that overpredicts event load. Metering accuracy does not solve those counterfactual incentives. A credible program needs a baseline method, validation procedure, event rule, audit trail, and treatment of missing data.

7 The reference implementation

7.1 Why exhaustive discrete clearing

Production markets use sophisticated optimization, network models, security analysis, and settlement systems. The reference implementation has a different objective: make every assumption inspectable. It uses only the Python standard library and clears small schedules by enumerating power levels on a declared quantum.

For each generator and interval, the candidate levels run from minimum to maximum in increments of the market quantum. A generator outage fixes output to zero. For each workload and interval, candidates run from its declared minimum to maximum. When telemetry is required but unavailable, the workload is fixed to fallback power.

For every joint candidate, the solver:

1. checks workload ramps, cumulative deadline energy, and restart events;
2. procures upward reserve from eligible generator headroom in ascending offer order;
3. checks total active-power balance;
4. finds connected network islands after line outages;
5. checks balance in every island;
6. solves the reduced linear angle equations by deterministic Gaussian elimination;
7. rejects thermal line overloads;
8. calculates submitted and audited true surplus;
9. retains the feasible candidate with highest submitted surplus.

Ties retain the first lexicographically generated schedule. A hard candidate-space limit prevents accidental use on a large production problem. This method is exponential. That is acceptable for the small adversarial fixtures and unacceptable for real system operations.

7.2 Generic data structures

The implementation defines nodes, lines, generators, fixed demands, workload bids, reserve requirements, clearing prices, and baseline claims. No acceptance test branches on a paper title, scenario name, or expected answer. Workload kind is descriptive metadata. Feasibility comes from numeric constraints shared by every kind.

This detail prevents a common research-code failure. A test-specific branch can make a named fixture pass without implementing the claimed abstraction. Here, congestion, outages, deadlines, and telemetry all flow through generic constraints.

7.3 Independent dispatch checks

The checker does not trust the residuals stored by the optimizer. It recomputes:

- nodal injection minus net line outflow;
- total generation minus total fixed and flexible demand;
- thermal line overload;
- reserve shortfall;
- deadline-energy shortfall;
- workload ramp violation.

A dispatch passes only when the maximum residual is within a declared tolerance. The settlement routine separately recomputes energy and total cash identity residuals.

7.4 What the software does not do

The solver does not calculate production LMPs. It does not solve AC power flow, unit commitment, security-constrained economic dispatch, or dynamic simulation. It has one upward-reserve product and no zonal deliverability. Generator ramps and start costs are omitted. Workload values are linear. Uncertainty is represented by scenarios supplied outside the solver, not by a stochastic program.

8 Computational scenarios

All scenarios in this section are COMPUTATIONAL and synthetic. They are regression fixtures, not observations of ERCOT, CAISO, a named data center, or a mining site.

8.1 Energy and reserve co-constraint

The first fixture has one node, one generator, 5 MW of fixed demand, and one AI bid. The generator can produce 20 MW, offer up to 5 MW of reserve, and has energy cost 20 per MWh. The AI workload bids 100 per MWh for up to 10 MW. The reserve requirement is 3 MW and the reserve offer is 2 per MWh of reserved capability.

The reserve constraint does not reduce AI load because the generator has enough unused capability. It does constrain the feasible generator state. If load rose to use the full 20 MW capability, the 3 MW reserve requirement would fail.

At an illustrative nodal energy price of 40 per MWh, total load pays 600 and the generator receives 600. Reserve credit and reserve charge are both 6. Both the energy identity and combined identity close to test tolerance.

Table 2: Single-node energy and reserve fixture

Quantity	Cleared result
Fixed demand	5 MW
AI load	10 MW
Generator energy	15 MW
Reserve award	3 MW
Generator headroom after energy	5 MW
Submitted surplus	694 currency units
Maximum checked balance residual	below 10^{-9} MW

8.2 Congestion allocates scarce transfer capability

The second fixture has west and east nodes joined by one line. A 10 MW generator at west costs 10 per MWh. A proof-of-work workload at west bids 80 per MWh. An AI inference workload at east bids 100 per MWh. Both can consume up to 10 MW.

With a 10 MW line, all available generation serves the higher bid at east. AI receives 10 MW, mining receives zero, and line flow is 10 MW. With a 5 MW line, only 5 MW can reach east. The remaining 5 MW serves the mining bid at west.

Table 3: Congestion changes allocation

Line capacity	AI at east	Mining at west	Submitted surplus
10 MW	10 MW	0 MW	900
5 MW	5 MW	5 MW	800

The result is not a universal priority rule. It follows from the submitted bids, one generator, and one line in this fixture. Reverse the bids and the allocation reverses. Add local generation and both loads may run.

8.3 Nodal congestion-rent identity

A related fixture has 5 MW of fixed demand at east and a 5 MW mining load at west. The west generator produces 10 MW, and 5 MW crosses the line east. Set the illustrative west price to 20 and east price to 50.

The west mining load pays $5 \times 20 = 100$. East fixed load pays $5 \times 50 = 250$. The west generator receives $10 \times 20 = 200$. The line rent is

$$(50 - 20) \times 5 = 150.$$

Thus $350 - 200 = 150$, exactly as Theorem 5.4 requires. The example demonstrates accounting closure at declared prices. It does not show how those two prices were formed.

8.4 Line outage and island balance

Consider 10 MW of fixed demand at east, a 10 MW cheap generator at west, and a 10 MW expensive generator at east. With the line available, west generation serves east demand. When the line is out, the two nodes form separate islands. The west generator falls to zero and the east generator supplies the full 10 MW.

Table 4: Line-outage redispatch

State	West generation	East generation	Line flow
Available	10 MW	0 MW	10 MW
Outage	0 MW	10 MW	0 MW

The island solver requires each connected component to balance. It does not allow an isolated deficit to be canceled by surplus in an electrically separate island.

8.5 Generator outage and reserve infeasibility

A one-node fixture has two 10 MW generators, 10 MW fixed demand, and a 5 MW reserve requirement. With both units available, one can provide energy while the other preserves reserve headroom. If the lower-cost unit is out, the remaining 10 MW unit must use all capability for energy. No upward headroom remains, so the instance is infeasible.

This is a useful negative test. A solver that checks only energy balance would accept the outage schedule. The reserve constraint correctly rejects it. The test still does not establish frequency response or reserve deliverability.

8.6 AI deadline, ramp, and restart

A two-interval AI batch uses zero-based interval labels 0 and 1. It needs 15 MWh by the end of the second interval, which is numbered 1. It begins at zero, can ramp up by at most 5 MW per interval, can reach 10 MW, and incurs a restart cost of 25. Each interval lasts one hour. The only feasible value-maximizing path is 5 MW in interval 0 and 10 MW in interval 1.

At a bid of 100 per MWh and zero generation cost, gross submitted value is 1,500. The one restart subtracts 25, leaving submitted surplus 1,475. A model that treats the load as instantaneously flexible would incorrectly allow 0 MW followed by 15 MW, even though the declared maximum and ramp both rule that out.

8.7 Telemetry failure

A 10 MW flexible load shares a one-node system with a 10 MW generator and a 5 MW reserve requirement. When telemetry is available, the market schedules the load at 5 MW and leaves 5 MW of headroom for reserve. When telemetry fails, the declared physical fallback is 10 MW. The full generator capability is then consumed by energy and the reserve requirement cannot be met. The problem becomes infeasible.

This scenario makes no claim that every telemetry failure produces a 10 MW fallback. The point is structural: the fallback must be declared. Treating telemetry loss as automatic zero load would manufacture flexibility that the operator cannot observe.

8.8 Strategic bid

One generator can serve exactly one 10 MW workload. The AI workload has a submitted and audited value of 100 per MWh. The mining workload has an audited value of 80 per MWh. Under truthful bids, AI receives 10 MW and audited true surplus is 1,000.

Now let mining submit 120 per MWh while its audited value remains 80. Mining wins the entire 10 MW. Submitted surplus rises to 1,200, but audited true surplus falls to 800. The optimizer did what it was asked to do. It maximized submitted bids. The scenario rejects any claim that the allocation is strategy-proof.

8.9 Baseline inflation

The baseline fixture in Section 6 reports a 4 MWh reduction and a payment of 400 even though the verified counterfactual implies no reduction. This test does not accuse a real participant of gaming. It verifies that the accounting code exposes payment attributable solely to baseline difference.

9 Adversarial market-design analysis

9.1 Baseline manipulation

A baseline program creates value from the difference between an estimated counterfactual and measured load. The participant may be able to influence both. Defenses can include matched control groups, day selection rules, weather adjustment, caps, exclusions, random audits, and penalties. Each defense changes error and participation incentives. No method can be assessed without the data-generating process and event frequency.

The market-design question is not merely whether the baseline formula is documented. It is whether a participant can profit by changing non-event load, withholding flexibility from baseline periods, shifting consumption across meters, or exploiting missing data. Tests should compare reported response with a counterfactual estimate that was not chosen after observing the desired payment.

9.2 Telemetry failure and stale control

Telemetry failure can produce at least four different states:

1. the load continues at its previous power;
2. a local controller moves it to a safe fallback;
3. the facility trips;
4. the operator receives stale or misleading data while the physical load follows another state.

Only the second and third states may look like curtailment, and neither should be assumed. A market qualification process must specify data quality, heartbeat, fallback, loss-of-signal behavior, remote-control authority, and recovery. Settlement needs rules for estimated or missing meter data.

9.3 Strategic bids and market power

Large flexible loads may be price takers in a broad market and still affect a constrained local node. A load can also own generation, storage, transmission rights, or related financial positions. Submitted willingness to pay need not equal social value. A market-power analysis should examine ownership, concentration, pivotality, congestion, and affiliated positions.

The reference model records both submitted and audited value precisely so the two concepts cannot be silently merged. It does not solve the mechanism-design problem.

9.4 Congestion and rebound

Low-price energy can attract new compute load. If many operators colocate behind the same constrained interface, congestion can erase the original price advantage and trigger network upgrades. Efficiency improvements can lower compute cost and expand total demand. A static dispatch captures the current line limit, but not the long-run entry response or induced investment.

Claims that compute “uses stranded energy” need an interval-level counterfactual. The energy may be curtailed in one season and scarce in another. A new durable load can also change generation retirement, transmission investment, and capacity procurement. Those effects require a dynamic market and planning model.

9.5 Outages and common-mode risk

The simulator models explicit line and generator outages. It does not model a common-mode event that disables a grid control channel, a data-center network, and a workload scheduler at once. Nor does it estimate correlated weather, fuel, cyber, or cooling failures. A project offering reliability service should identify dependencies shared by the grid and the load.

9.6 Reserve double counting

The same megawatt cannot simultaneously be sold as multiple fully overlapping services unless the tariff and physical capability permit stacking. A compute site might advertise load curtailment, battery discharge, backup generation, and reserve from the same interconnection margin. The accounting must show which resource, meter, interval, and activation path supports each award.

The reference model avoids double counting generator headroom within its one reserve product through eq. (5). A multi-product extension would need coupled capability constraints.

10 What a production study must add

10.1 Alternating-current feasibility

A production interconnection study needs bus voltage, reactive-power capability, transformer limits, losses, contingencies, protection, and possibly unbalanced distribution analysis. Large power-electronic loads can also raise harmonic, power-quality, and control-interaction questions. None is visible in eq. (1).

10.2 Dynamic stability

Frequency and voltage stability are time-domain properties. A valid study may need generator and inverter controls, load response, fault behavior, protection timing, under-frequency action, and recovery. The relevant contingencies and performance criteria must come from the applicable reliability and interconnection rules.

The DC model is incapable of proving:

- frequency nadir or rate of change of frequency;
- voltage recovery after a fault;
- small-signal or transient stability;
- fault current and protection coordination;
- ride-through performance;
- reserve response time or sustained response.

10.3 Security-constrained and unit-commitment detail

Production dispatch often includes unit commitment, generator ramping, minimum run times, start costs, security constraints, reserves by class, and contingency deliverability. Flexible loads may have nonconvex block bids or minimum run sizes. The discrete reference model can express some nonconvex workload schedules but omits most generator commitment detail.

10.4 Distribution and retail terms

A wholesale nodal price may not be the bill faced by a data center or mine. Retail tariffs can include demand charges, coincident-peak charges, fixed charges, power-factor penalties, riders, taxes, and negotiated service terms. Transmission and distribution upgrade costs can dominate short-run energy price differences. A site study must reconcile wholesale dispatch value with the actual customer bill and cost allocation.

10.5 Emissions and water

Low price is not synonymous with low emissions. Marginal generation, local air pollution, water use, backup generation, and embodied infrastructure can change the social account. If a flexible load moves between intervals or regions, the analysis should use marginal consequences for the same counterfactual and time horizon. Average annual grid mix is not enough.

11 Market design for responsive compute

11.1 Participation product

A credible market product should define what the load sells. Possibilities include scheduled energy consumption, price-responsive demand, upward curtailment, regulation, contingency reserve, emergency service, or distribution congestion relief. Each product needs:

- activation trigger and response deadline;
- minimum and maximum response;
- telemetry and meter resolution;
- baseline or direct-control measurement;
- sustainment and recovery rules;
- location and deliverability;
- nonperformance consequences;
- interaction with other awards.

Describing the product as “flexibility” is not enough to settle it.

11.2 Direct scheduling versus baseline response

Direct scheduling can avoid part of the baseline problem. If a controllable load submits a consumption bid and is dispatched to an explicit megawatt level, performance can be compared with that schedule and meter data. This still requires telemetry, deviation settlement, and a rule for outages. It does not require estimating the same counterfactual as a pure load-reduction product.

Baseline response can be appropriate when direct scheduling is unavailable, but it should carry an uncertainty and gaming account. The expected benefit of the program should be net of baseline error, administrative cost, and induced behavior.

11.3 AI queue design

An AI operator can expose flexibility without revealing individual customer content. A market-facing scheduler can aggregate workload envelopes:

- must-run power for latency-critical service;
- deferrable energy with deadlines;
- interruptible jobs with checkpoint states;
- restart cost and recovery time;
- maximum ramp by interval;
- fallback power under lost control.

The market needs the envelope, not private task payloads. The operator remains responsible for making the envelope conservative enough to deliver.

11.4 Mining fleet design

A mining operator can aggregate machine groups by efficiency, cooling zone, firmware control, and restart behavior. The site bid can then represent a piecewise demand curve instead of

an all-or-nothing block. Pool and hedge positions should be reflected in private value, while public claims about settlement assurance require a separate attribution analysis.

11.5 Scarcity and emergency operation

Price response and emergency control are different authorities. A high price may induce voluntary curtailment. An emergency program can require performance under specified conditions. Interconnection agreements may permit protective trips. These paths have different notice, compensation, testing, and liability.

A market design should not count voluntary price response as firm emergency capacity without a binding qualification rule.

12 Empirical research program

12.1 Minimum dataset

An empirical study should publish one row per resource, node, and interval, with stable identifiers and the following fields:

- interval start, duration, and time zone;
- scheduled, metered, baseline, and fallback power;
- bid curve and cleared quantity;
- task class, deadline energy, and service-loss rule;
- ramp limits, restart events, and recovery time;
- telemetry availability and data-quality flags;
- nodal energy and congestion price components;
- reserve product, award, activation, and performance;
- line or zone constraints relevant to deliverability;
- generator or grid emissions counterfactual, if evaluated;
- outage and event labels;
- settlement adjustments, uplift, and penalties.

12.2 Causal questions

Several distinct causal questions are often compressed into one headline:

1. Did a price or dispatch instruction change compute load?
2. Did that change reduce system production cost?
3. Did it relieve the binding local constraint?
4. Did it preserve or improve reliability under the relevant contingency?
5. Did it reduce emissions or merely shift them?
6. Did induced entry increase long-run energy use?
7. Did the resulting intelligence or settlement service create realized economic value?

Each question needs its own counterfactual. A metered reduction can answer the first only if the baseline is credible. It does not by itself answer the other six.

12.3 Identification designs

Randomized dispatch, where operationally safe, can estimate load response and baseline error. Price discontinuities or event thresholds may support quasi-experimental analysis. Matched control groups can help when treatment and control sites share weather and operational patterns. Structural estimation can connect bids, schedules, and latent service costs.

Every design must account for anticipation. A participant that knows baseline days or event rules may alter load before treatment. Frequent dispatch can also contaminate historical baselines. Missing telemetry is rarely random.

12.4 Outcome reporting

Results should separate:

- gross energy payment;
- congestion and reserve payment;
- baseline adjustment;
- nonperformance penalty;
- workload service loss;
- restart and recovery cost;
- grid production-cost change;
- upgrade and interconnection cost;
- emissions and other external effects.

Reporting only the participant’s electricity savings does not identify system value.

13 Verification contract

13.1 Executable checks

The reference suite contains eleven generic tests:

Table 5: Executable verification matrix

Test	Acceptance condition
Energy and reserve clearing	workload, generation, and reserve match the declared optimum; independent residual check passes
Single-node settlement	energy payments equal energy credits; reserve charges equal reserve credits
Congestion allocation	a binding 5 MW line splits scarce generation across the two workload nodes
Congestion rent	nodal load payment minus generator credit equals line rent
Line outage	island balance forces local generation
Generator outage	energy-only feasibility is rejected when reserve headroom is absent

Test	Acceptance condition
AI deadline and ramp	cumulative delivery, ramp, and restart cost all bind as declared
Telemetry failure	physical fallback is used and reserve infeasibility is reported
Baseline inflation	excess payment is exposed when claimed and verified counterfactuals differ
Strategic bid	submitted-surplus allocation can lower audited true surplus
Input validation	undeclared nodes and excessive candidate spaces are rejected

13.2 Reproduction

From the repository root, the focused computational evidence is reproduced with:

```
pytest -q tests/test_energy_markets.py
```

The paper is compiled in an isolated temporary directory by:

```
scripts/build_paper.sh \
  energy-markets-for-machine-intelligence-and-cryptographic-settlement
```

The build gate requires the final PDF, no unresolved references or citations, and no LaTeX error. The code and tests use no network service and no random seed.

13.3 Interpretation of a passing suite

A passing suite establishes that the implementation behaves as declared on the tested finite instances. It does not validate the synthetic input values. It does not demonstrate scalability. It does not certify a real grid or market participant. It does not convert a DC dispatch into a dynamic-stability result.

14 Limits and failure conditions

Table 6: Model limits and required next evidence

Omission	Consequence	Required extension
AC voltage and reactive power	DC feasibility can accept an electrically unacceptable schedule	AC power flow and contingency study

Omission	Consequence	Required extension
Frequency dynamics	reserve headroom is not response	dynamic models, controls, event tests
Network losses	line-rent identity omits marginal-loss surplus	loss-aware power flow and settlement account
Unit commitment	generator nonconvex costs are understated	commitment, ramp, start, and minimum-run model
Multiple reserve products	service capabilities can be double counted	coupled product and deliverability constraints
Uncertainty	deterministic schedule can be brittle	scenarios, recourse, and probabilistic adequacy
Endogenous prices	declared settlement prices may not support dispatch	continuous or mixed-integer market-pricing method
Market power	bid surplus can diverge from true surplus	ownership, pivotality, and conduct analysis
Baseline identification	measured load does not reveal counterfactual	validated causal measurement design
Long-run entry	static dispatch misses rebound and investment	dynamic equilibrium and planning model
Retail tariffs	wholesale price is not customer bill	tariff and contract reconciliation
External effects	private bid is not social value	matched emissions, water, land, and public-cost account

15 Decision protocol for a real project

A practical evaluation can use the following sequence.

1. **Resolve the boundary.** Name the meter, node, interval, counterfactual, market product, and responsible entity.
2. **Measure the feasible set.** Record minimum and maximum import, ramp, deadline, restart, fallback, and control latency.
3. **Check qualification.** Confirm the current interconnection, scheduling, metering, telemetry, and certification requirements.
4. **Clear energy and reserve together.** Do not spend headroom twice.
5. **Test congestion and outages.** Include line and generator contingencies and island balance.
6. **Audit settlement.** Close energy, congestion, reserve, uplift, and baseline accounts.
7. **Attack the mechanism.** Test baseline manipulation, stale telemetry, strategic bids, affiliated positions, and missing data.
8. **Run engineering studies.** Use AC and dynamic tools appropriate to the connection and claimed service.
9. **Estimate realized value.** Connect consumed energy to quality-adjusted intelligence or settlement assurance, then to realized economic outcomes.

10. **Report uncertainty and nonclaims.** Separate executable checks from empirical estimates and unresolved questions.

16 Claim ledger

Claim JS-C012. AI and proof-of-work workloads have different feasible flexibility sets once deadlines, ramps, restart costs, minimum service, and telemetry are declared. Status: COMPUTATIONAL in the reference fixtures; OPEN for any unnamed real site.

System active-power balance. Summed nodal balance equals total generation minus fixed and flexible demand. Status: PROVED under Assumption 5.1.

Lossless nodal settlement. Load payment minus generator credit equals line congestion rent at any complete nodal price vector. Status: PROVED by Theorem 5.4.

Reserve cash closure. Total cash closes when reserve charges are allocated to equal reserve credits. Status: PROVED as an accounting convention, not as a tariff claim.

Congestion allocation. A binding synthetic line changes allocation between AI and mining bids. Status: COMPUTATIONAL.

Telemetry consequence. A synthetic physical fallback can make an otherwise feasible reserve schedule infeasible. Status: COMPUTATIONAL.

Baseline overpayment. An inflated reported baseline can create payment without verified reduction. Status: COMPUTATIONAL for the declared arithmetic; real-world incidence is OPEN.

Strategic bidding. Submitted-surplus maximization can lower audited true surplus when bids are not truthful. Status: COMPUTATIONAL.

Claim JS-C013. DC dispatch feasibility does not prove dynamic stability. Status: OBSTRUCTED by model scope.

Automatic grid benefit. Flexible compute is necessarily grid positive. Status: OBSTRUCTED. Location, timing, baseline, deliverability, rebound, and omitted engineering can reverse the result.

17 Conclusion

Machine intelligence and cryptographic settlement can participate in electricity markets only through physical loads, meters, control systems, bids, and rules. Their shared energy input does not erase their different deadlines, restart losses, curtailment ranges, and economic outputs.

The bounded result of this paper is useful. A transparent discrete market can co-optimize energy and a simple reserve product, enforce workload feasibility, solve lossless DC line flows, and expose congestion, outage, telemetry, baseline, and strategic-bid failure modes. Within that boundary, active-power balance and nodal settlement close exactly.

The boundary is equally important. No DC-OPF result proves voltage, frequency, or dynamic stability. No baseline payment proves a real counterfactual reduction. No submitted bid proves social value. No flexible upper bound proves a qualified market service. The next research and engineering step is therefore not a larger slogan. It is a site-specific feasible set,

current market rule, adversarial settlement audit, and the AC and dynamic studies needed for the service being claimed.

A Compact mathematical program

For a continuous relaxation without restart binaries, the core problem is

$$\begin{aligned} \max_{p,r,x,\theta,f} \quad & \Delta \sum_{w,t} \beta_w x_{wt} - \Delta \sum_{g,t} c_g p_{gt} - \Delta \sum_{g,t} k_g r_{gt} \\ \text{subject to} \quad & \text{eqs. (1) to (10).} \end{aligned} \tag{22}$$

With linear bids and costs, this relaxation is a linear program. Restart indicators, block bids, and discrete power quanta make the reference problem nonconvex. The exhaustive solver handles those small nonconvex instances directly.

B Settlement derivation by incidence matrix

Let A be the node-line incidence matrix with +1 at each line's from-node and -1 at its to-node. Let net injection be $y_t = g_t - d_t - x_t$. Nodal balance is

$$y_t = Af_t.$$

For nodal price vector λ_t ,

$$P_t^D - C_t^G = -\Delta \lambda_t^\top y_t = -\Delta \lambda_t^\top Af_t.$$

For line ℓ , the corresponding component of $-A^\top \lambda_t$ is $\lambda_{d(\ell),t} - \lambda_{o(\ell),t}$. Therefore

$$P_t^D - C_t^G = \Delta \sum_{\ell} (\lambda_{d(\ell),t} - \lambda_{o(\ell),t}) f_{\ell t}.$$

Summing over intervals reproduces eq. (17).

C Audit schema

Table 7: Minimum machine-readable fields

Field	Unit or type	Meaning
node_id	text	electrical settlement or modeled node
interval	integer and times-tamp	ordered dispatch interval with time zone
interval_hours	hours	conversion from MW to MWh
line_capacity	MW	directional absolute thermal limit

Field	Unit or type	Meaning
line_availability	Boolean	interval outage state
generator_minimum	MW	available minimum output
generator_maximum	MW	available maximum output
energy_offer	currency per MWh	marginal energy offer or cost assumption
reserve_maximum	MW	offered upward reserve bound
reserve_offer	currency per MW	reserve opportunity-cost assumption
fixed_demand	MW	demand outside flexible bids
workload_kind	controlled vocabulary	descriptive class, not a solver branch
workload_minimum	MW	must-run or contractual floor
workload_maximum	MW	facility and service ceiling
workload_bid	currency per MWh	submitted dispatch value
audited_value	currency per MWh	ex-post welfare value assumption
initial_power	MW	state before first interval
ramp_up and ramp_down deadline	MW per interval interval or null	intertemporal movement limits last interval for required cumulative energy
required_energy	MWh	minimum delivery through deadline
restart_cost	currency	cost of zero-to-positive transition
telemetry_available	Boolean	modeled control eligibility state
fallback_power	MW	physical schedule under telemetry loss
reported_baseline	MW	baseline used for claimed response
verified_counterfactual	MW	audit estimate of no-event load
metered_power	MW	observed interval consumption

D Reviewer checklist

1. Are node, interval, meter, and counterfactual explicit?
2. Are AI and mining feasible sets measured rather than inferred from labels?
3. Does the reserve account preserve headroom after energy dispatch?
4. Are outage cases balanced by electrical island?
5. Is telemetry failure assigned a physical fallback?
6. Is baseline response separated from metered load?
7. Can strategic bids be distinguished from audited value?
8. Does settlement close across energy, congestion, and reserve?
9. Are marginal losses omitted, modeled, or separately allocated?
10. Is the result incorrectly described as a frequency or voltage proof?
11. Have current operator rules and interconnection requirements been checked?
12. Are all scenario inputs labeled observed, estimated, synthetic, open, or obstructed?

References

- [1] Federal Energy Regulatory Commission. *Energy Markets Primer*. January 2024. https://www.ferc.gov/sites/default/files/2024-01/24_Energy-Markets-Primer_0117_DIGITAL_0.pdf.
- [2] Federal Energy Regulatory Commission. Understanding energy markets. Accessed 23 July 2026. <https://www.ferc.gov/understanding-energy-markets>.
- [3] Federal Energy Regulatory Commission. Reports on demand response and advanced metering. Updated 12 December 2025. <https://www.ferc.gov/power-sales-and-markets/demand-response/reports-demand-response-and-advanced-metering>.
- [4] United States Department of Energy. *Recommendations on Powering Artificial Intelligence and Data Center Infrastructure*. July 2024. https://www.energy.gov/sites/default/files/2024-08/Powering%20AI%20and%20Data%20Center%20Infrastructure%20Recommendations%20July%202024_1.pdf.
- [5] Electric Reliability Council of Texas. Large load integration. Accessed 23 July 2026. <https://www.ercot.com/services/rq/large-load-integration>.
- [6] Electric Reliability Council of Texas. Current Planning Guide. Accessed 23 July 2026. <https://www.ercot.com/mktrules/guides/planning/current>.
- [7] Electric Reliability Council of Texas. *ERCOT Nodal Protocols*. 1 July 2026. <https://www.ercot.com/files/docs/2026/06/18/July-1-2026-Nodal-Protocols.pdf>.
- [8] Electric Reliability Council of Texas. Load resource participation in the ERCOT markets. Accessed 23 July 2026. <https://www.ercot.com/services/programs/load/laar>.
- [9] California Independent System Operator. Demand response. Accessed 23 July 2026. <https://www.caiso.com/generation-transmission/load/demand-response>.
- [10] California Independent System Operator. Metering and telemetry. Accessed 23 July 2026. <https://www.caiso.com/generation-transmission/metering-telemetry>.
- [11] California Independent System Operator. *Business Practice Manual for Demand Response*. Current document library accessed 23 July 2026. <https://bpmcm.caiso.com/Pages/BPMDetails.aspx?BPM=Demand%20Response>.
- [12] North American Electric Reliability Corporation. Standards. Accessed 23 July 2026. <https://www.nerc.com/pa/Stand/Pages/default.aspx>.
- [13] North American Electric Reliability Corporation. *Glossary of Terms Used in NERC Reliability Standards*. Current glossary accessed 23 July 2026. https://www.nerc.com/globalassets/standards/reliability-standards/glossary_of_terms.pdf.

- [14] F. C. Schweppe, M. C. Caramanis, R. D. Tabors, and R. E. Bohn. *Spot Pricing of Electricity*. Kluwer Academic Publishers, 1988. <https://doi.org/10.1007/978-1-4613-1683-1>.
- [15] A. J. Wood, B. F. Wollenberg, and G. B. Sheble. *Power Generation, Operation, and Control*. Third edition, Wiley, 2013.
- [16] R. D. Zimmerman, C. E. Murillo-Sanchez, and R. J. Thomas. MATPOWER: steady-state operations, planning, and analysis tools for power systems research and education. *IEEE Transactions on Power Systems*, 26(1):12 to 19, 2011. <https://doi.org/10.1109/TPWRS.2010.2051168>.
- [17] MATPOWER. Mathematical model object and DC optimal power flow documentation. Version 8.1 documentation, accessed 23 July 2026. <https://matpower.org/doc/dev-manual/math-model.html>.
- [18] Z. Liu, M. Lin, A. Wierman, S. H. Low, and L. L. H. Andrew. Greening geographical load balancing. *ACM SIGMETRICS Performance Evaluation Review*, 39(1):193 to 204, 2011. <https://doi.org/10.1145/2007116.2007139>.
- [19] Z. Liu, I. Liu, S. H. Low, and A. Wierman. Pricing data center demand response. *ACM SIGMETRICS Performance Evaluation Review*, 42(1):111 to 123, 2014. <https://doi.org/10.1145/2591971.2592004>.
- [20] A. Wierman, Z. Liu, I. Liu, and H. Mohsenian-Rad. Opportunities and challenges for data center demand response. In *2014 International Green Computing Conference*, 2014. <https://doi.org/10.1109/IGCC.2014.7039172>.
- [21] F. Yang and A. A. Chien. Large-scale and extreme-scale computing with stranded green power: opportunities and costs. *IEEE Transactions on Parallel and Distributed Systems*, 29(5):1103 to 1116, 2018. <https://doi.org/10.1109/TPDS.2017.2782677>.
- [22] A. Menati, K. Lee, and L. Xie. Modeling and analysis of utilizing cryptocurrency mining for demand flexibility in electric energy systems: a synthetic Texas grid case study. *IEEE Transactions on Energy Markets, Policy and Regulation*, 1(1):1 to 10, 2023. <https://doi.org/10.1109/TEMPR.2022.3230953>.
- [23] A. Menati, X. Zheng, K. Lee, R. Shi, P. Du, C. Singh, and L. Xie. High resolution modeling and analysis of cryptocurrency mining’s impact on power grids: carbon footprint, reliability, and electricity price. *Advances in Applied Energy*, 10:100136, 2023. <https://doi.org/10.1016/j.adapen.2023.100136>.