

GPU-Accelerated Nonlinear Multi-Period AC Optimal Power Flow for Large-Scale Power-Hydrogen Systems

Geunseo Song^{ab}, Dirk Lauinger^{bc}, Sungho Shin^{b*}, and Jonggeol Na^{a*}

^a Ewha Womans University, Department of Chemical Engineering and Materials Science, Seoul, Republic of Korea

^b Massachusetts Institute of Technology, Department of Chemical Engineering, Cambridge, Massachusetts, USA

^c Massachusetts Institute of Technology, MIT Energy Initiative, Cambridge, Massachusetts, USA

* Corresponding Authors: sushin@mit.edu (S. Shin), jgna@ewha.ac.kr (J. Na)

ABSTRACT

The growing penetration of renewable energy sources and power-to-hydrogen (P2H) systems demands high-fidelity, large-scale optimization frameworks that capture the nonlinear physics of both AC power flow and hydrogen thermodynamics. However, existing approaches rely on DC approximations and simplified electrolyzer models, neglecting critical operational constraints. As a result, accurately modeling such systems leads to large-scale nonlinear programs that are computationally intractable for conventional CPU-based solvers. This motivates the need for scalable optimization frameworks capable of handling both physical fidelity and computational complexity. This paper proposes a fully GPU-native framework for solving large-scale multi-period AC optimal power flow (AC-OPF) problems with integrated power-to-hydrogen systems. High-fidelity thermodynamic models of hydrogen production, compression, cooling, and storage are coupled with AC power flow constraints, resulting in large-scale nonlinear programs (NLPs) with up to 14.4 million variables and 22.4 million constraints in the largest benchmark case (9, 591-bus network with a 168-hour horizon). To enable scalable solutions, condensed-space KKT reformulations and GPU-accelerated sparse Cholesky factorization are employed within an end-to-end GPU optimization pipeline. Numerical results on benchmark networks up to 9, 591 buses demonstrate 10–600× speedups over CPU solvers, whereas CPU-based solvers fail to converge on the largest instances. Operational studies further highlight the importance of thermodynamic constraints in realistic hydrogen system scheduling.

Keywords: AC optimal power flow, power-to-hydrogen, integrated energy systems, GPU-accelerated optimization, nonlinear programming

1. INTRODUCTION

Global electricity demand is growing rapidly due to the electrification of industrial processes, the expansion of data centers, and the increasing penetration of electric transportation systems [1]. Achieving carbon neutrality under these trends requires a fundamental transformation of power systems toward renewable energy sources (RES). However, the inherent intermittency and uncertainty of RES introduce significant temporal mismatches between electricity supply and demand, posing critical challenges to secure and efficient system operation.

Addressing these challenges necessitates flexible energy storage solutions that enable sector coupling

across multiple energy domains. Among the available options, Power-to-Hydrogen (P2H) technology offers a promising solution by converting surplus electricity into hydrogen, enabling long-duration storage and cross-sectoral decarbonization [2].

Despite the potential of integrated systems, existing optimization frameworks often rely on significant simplifications to manage computational complexity. On the grid side, many studies utilize Direct Current (DC) power flow approximations, which ignore voltage limits and reactive power requirements [3]. On the hydrogen side, electrolyzers are frequently represented as idealized, controllable electrical loads [4]. Such formulations neglect critical physical characteristics of hydrogen systems, including the nonlinear auxiliary power

consumption of balance-of-plant components (e.g., compressors and cooling systems) and the thermodynamic dynamics of hydrogen storage vessels, such as pressure and temperature evolution. Ignoring these effects can lead to inaccurate assessments of system flexibility and suboptimal operational decisions.

To address these limitations, this paper proposes a high-fidelity, integrated optimization framework that couples a multi-period Alternating Current Optimal Power Flow (AC-OPF) model with detailed hydrogen production and storage dynamics. The AC-OPF formulation explicitly captures voltage magnitudes, active and reactive power flows, and transmission losses, while the hydrogen subsystem is modeled with rigorous operational and thermodynamic constraints. This integrated formulation enables a physically consistent representation of the coupled power-hydrogen system across multiple time periods.

The resulting optimization problem is a large-scale, nonlinear program (NLP) characterized by strong temporal coupling and complex network constraints. Solving such problems using conventional CPU-based nonlinear solvers becomes computationally prohibitive for realistically sized systems, particularly at the national scale.

To overcome this computational challenge, the proposed framework is implemented using ExaModels.jl, a GPU-based algebraic modeling system, and solved with MadNLP.jl, a GPU-accelerated nonlinear solver [5]. By leveraging the massive parallelism of modern graphics processing units (GPUs), this computational strategy enables the efficient solution of large-scale, high-fidelity integrated energy system models that were previously intractable.

The proposed approach is validated through case studies ranging from regional to national-scale power networks. The results demonstrate both the operational importance of detailed hydrogen process modeling and the substantial computational advantages of GPU-accelerated nonlinear optimization.

The remainder of this paper is organized as follows. Section 2 presents the mathematical formulation of the multi-period AC-OPF and hydrogen system models. Section 3 describes the GPU-accelerated computational framework. Section 4 reports numerical results and performance analyses, and Section 5 concludes the paper.

2. MATHEMATICAL MODELING

This section presents the mathematical formulation of the proposed multi-period AC-OPF problem with integrated hydrogen production and storage. Figure 1 illustrates the overall structure of the coupled power-hydrogen system, including the electrical network, electrolyzer, compressor, cooling system, and hydrogen storage tank, as well as their physical and operational interactions.

The proposed formulation extends the conventional

AC power flow representation [6] by explicitly incorporating the inter-temporal dynamics and operational constraints of hydrogen subsystems. The model is implemented using ExaModelsPower.jl, an open-source modeling library for creating GPU-compatible nonlinear AC optimal power flow models [7].

The integrated system is formulated as a large-scale NLP problem over a discrete time horizon $t \in \mathcal{T}$. Building upon [8], the proposed formulation augments the basic power-to-hydrogen model with rigorous thermodynamic constraints for compression and cooling. At selected buses $i \in \mathcal{H} \subseteq \mathcal{N}$, hydrogen production and storage facilities are installed and electrically coupled to the power grid. Each hydrogen facility consists of a proton exchange membrane (PEM) electrolyzer, a compressor for pressurized hydrogen storage, and a cooling system to maintain safe operating temperatures. The coupling between the power grid and the hydrogen subsystem is realized through the electrical power consumption of these components, which directly enters the nodal power balance equations of the grid.

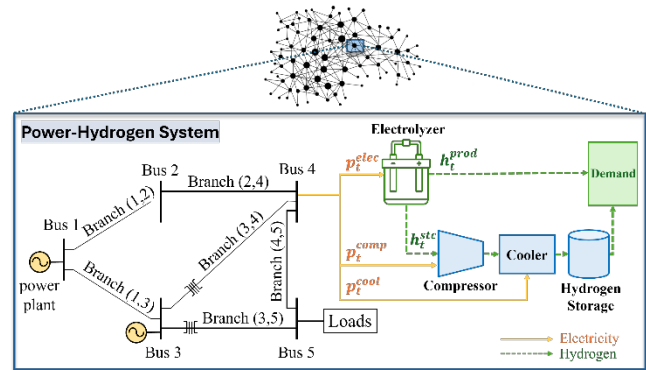


Figure 1. Schematic of the integrated power-hydrogen system. Selected buses in the transmission network are equipped with hydrogen production facilities, each comprising a PEM electrolyzer, compressor, and cooler. The electrical power consumption of these components enters the nodal power balance equations of the grid.

2.1 Multi-period AC-OPF with Integrated Hydrogen Loads

The core of the integrated energy system is the nonlinear AC power flow model, which ensures the physical feasibility of power system operation. For each bus $i \in \mathcal{N}$ and time period $t \in \mathcal{T}$, the nodal active and reactive power balance equations are formulated as

$$\sum_{g \in \mathcal{G}_i} p_{i,t}^g - p_{i,t}^d - p_{i,t}^{H2} = \sum_{j \in \mathcal{N}_i} p_{ij,t}, \forall i \in \mathcal{N}, t \in \mathcal{T} \quad (1)$$

$$\sum_{g \in \mathcal{G}_i} q_{i,t}^g - q_{i,t}^d = \sum_{j \in \mathcal{N}_i} q_{ij,t}, \forall i \in \mathcal{N}, t \in \mathcal{T} \quad (2)$$

where $p_{i,t}^g$ and $q_{i,t}^g$ denote the active and reactive power

generation, and $p_{i,t}^d$ and $q_{i,t}^d$ represent the fixed active and reactive power demands, respectively. The term $p_{i,t}^{H2}$ denotes the total active power consumption of the hydrogen facility. This term is non-zero only for the subset of buses equipped with hydrogen systems, $i \in \mathcal{H} \subseteq \mathcal{N}$, and is decomposed as

$$p_{i,t}^{H2} = p_{i,t}^{elec} + p_{i,t}^{comp} + p_{i,t}^{cool}, \quad (3)$$

where $p_{i,t}^{elec}$, $p_{i,t}^{comp}$, and $p_{i,t}^{cool}$ denote the electrical power consumption of the PEM electrolyzer, hydrogen compressor, and cooling system, respectively.

The branch active and reactive power flows between buses i and j are defined using the standard polar form of the AC power flow equations as

$$p_{ij,t} = G_{ij}V_{i,t}^2 - V_{i,t}V_{j,t}(G_{ij}\cos\theta_{ij,t} + B_{ij}\sin\theta_{ij,t}), \quad (4)$$

$$q_{ij,t} = -B_{ij}V_{i,t}^2 - V_{i,t}V_{j,t}(B_{ij}\cos\theta_{ij,t} + G_{ij}\sin\theta_{ij,t}), \quad (5)$$

where $V_{i,t}$ and $\theta_{i,t}$ are the voltage magnitude and angle at bus i , and G_{ij} and B_{ij} are the elements of the bus admittance matrix. Operational security is ensured by enforcing voltage magnitude limits, generator output limits, and branch flow limits.

2.2 Hydrogen Production Model

For notational simplicity, the bus index i is omitted in the hydrogen subsystem equations. All hydrogen-related variables are defined for each hydrogen-equipped bus $i \in \mathcal{H}$ and time period $t \in \mathcal{T}$, unless otherwise stated. The hydrogen production system is based on a PEM electrolyzer, modeled based on electrochemical principles described in [9]. The hydrogen generation rate is modeled as a function of the electrolyzer current. The hydrogen production rate h_t^{prod} is given by

$$h_t^{prod} = \frac{3600 \cdot N_c \cdot A \cdot M_{H_2} \cdot i_t \cdot \eta_F(i_t)}{2F}, \quad (6)$$

where N_c is the number of cells, A is the electrode active area, M_{H_2} is the molar mass of hydrogen, i_t is the current density, F is Faraday's constant, and $\eta_F(i_t)$ is the Faradaic efficiency.

The Faradaic efficiency is modeled as a logarithmic function of the current density following,

$$\eta_F(i_t) = 0.98 - 0.02\ln(i_t), \quad (7)$$

where the current density i_t is restricted to the practical operating range $0.5 \leq i_t \leq 2.5 \text{ A/cm}^2$.

The electrical power consumption of the electrolyzer is modeled as

$$p_t^{elec} = V_c(i_t) \cdot i_t \cdot A \cdot N_c, \quad (8)$$

where $V_c(i_t)$ is the voltage of a single cell, which depends nonlinearly on the operating current density. The electrochemical parameters for cell voltage are adopted

from [10].

2.3 Compressor and Cooling Power Models

Hydrogen produced by the electrolyzer is compressed prior to injection into the storage tank. The compressor power consumption is modeled based on isentropic compression:

$$p_t^{comp} = \frac{\gamma R T_{elec}}{(\gamma - 1)\eta_{comp}} \cdot n_t^{stc} \cdot \left[\left(\frac{P_t^{comp,out}}{P_{H_2}} \right)^{\frac{\gamma-1}{\gamma}} - 1 \right], \quad (9)$$

where γ is the specific heat ratio, R is the gas constant, T_{elec} is the electrolyzer outlet temperature, η_{comp} is the compressor efficiency, n_t^{stc} is the molar flow rate entering the storage tank, and $P_t^{comp,out}$ and P_{H_2} denote the compressor outlet and inlet pressures, respectively.

The outlet temperature of the compressor is given by

$$T_t^{comp,out} = T_{elec} + \frac{T_{elec}}{\eta_{comp}} \left[\left(\frac{P_t^{comp,out}}{P_{H_2}} \right)^{\frac{\gamma-1}{\gamma}} - 1 \right], \quad (10)$$

The cooling system removes excess thermal energy to maintain safe tank temperatures. The cooling power is derived from a simple energy balance:

$$p_t^{cool} = \frac{Q}{COP} = \frac{h_t^{stc} \cdot c_p \cdot \Delta T}{COP}, \quad (11)$$

where c_p is the specific heat capacity, ΔT is the required temperature reduction, and COP denotes the coefficient of performance of the cooling system.

2.4 Hydrogen Storage Dynamics and Thermodynamic Constraints

The hydrogen storage tank is modeled using dynamic mass and energy balance equations. The hydrogen inventory evolves as

$$h_t^{st} = h_{t-1}^{st} + h_t^{stc} - h_t^{std}, \quad (12)$$

where h_t^{stc} and h_t^{std} denote hydrogen charged into and discharged from the storage tank, respectively. The hydrogen inventory in the storage tank is bounded by its minimum and maximum storage capacities:

$$S^{min} \leq h_t^{st} \leq S^{max}, \quad (13)$$

The tank temperature dynamics are governed by the internal energy balance:

$$c_v \cdot n_t^{st} \cdot T_t^{tank} = c_v \cdot n_{t-1}^{st} \cdot T_{t-1}^{tank} + c_p \cdot n_t^{stc} \cdot T_{tank,in} - c_p \cdot n_t^{std} \cdot T_t^{tank}, \quad (14)$$

where c_v is the constant-volume specific heat, n_t^{st} is the molar quantity stored in the tank, and $T_{tank,in}$ denotes the inlet temperature after cooling.

The tank pressure is determined by the ideal gas law:

$$P_t^{tank} = \frac{n_t^{st} R T_t^{tank}}{V_{tank}}, \quad (15)$$

where V_{tank} denotes the storage volume.

To ensure operational safety and compatibility with downstream systems, the compressor outlet pressure satisfies

$$P_t^{comp,out} = P_t^{tank} + \Delta P_{margin}, \quad (16)$$

where ΔP_{margin} represents the pressure margin required for reliable injection into the storage vessel.

3. GPU-BASED COMPUTATIONAL FRAMEWORK

This section presents the computational framework developed to solve the proposed large-scale, multi-period nonlinear AC-OPF problem with integrated hydrogen production and storage. The framework adopts a fully GPU-native optimization pipeline based on ExaModels.jl for algebraic modeling, MadNLP.jl for nonlinear programming, and GPU-accelerated sparse linear solvers for Newton step computation. This end-to-end GPU execution enables scalable solution of tightly coupled, high-dimensional energy system optimization problems.

3.1 Computational Challenges in Large-Scale Nonlinear Optimal Power–Hydrogen Flow

The integrated power–hydrogen optimization problem formulated in Section 2 can be expressed as a large-scale NLP problem:

$$\begin{aligned} \min_x & f(x) \\ \text{s. t.} & g(x) = 0, \\ & h(x) \leq 0 \end{aligned}$$

where x comprises generator outputs, bus voltages, branch flows, and hydrogen system states (e.g., current density, tank pressure, temperature, and stored mass) across all buses and time periods.

Dimensionality and inter-temporal coupling: The problem dimension scales with both network size and time horizon length. The incorporation of hydrogen dynamics introduces additional state variables and constraints with strong inter-temporal coupling, substantially increasing memory requirements and computational complexity. For large-scale instances, this leads to NLPs with tens of millions of variables and constraints.

Nonlinearity and nonconvexity: The AC power flow equations introduce strong nonlinearities and nonconvexity, necessitating the use of Newton-type iterative methods to achieve reliable convergence. Each iteration requires solving a Karush–Kuhn–Tucker (KKT) system:

$$\begin{bmatrix} \nabla^2 L(x, \lambda) & J(x)^T \\ J(x) & 0 \end{bmatrix} \begin{bmatrix} \Delta x \\ \Delta \lambda \end{bmatrix} = - \begin{bmatrix} \nabla L(x, \lambda) \\ g(x) \end{bmatrix}$$

where $L(x, \lambda)$ is the Lagrangian and $J(x)$ is the Jacobian of constraints.

These systems are typically ill-conditioned due to wide variable scaling, tight operational constraints, and strong coupling between electrical and thermodynamic subsystems.

Large-scale sparse linear algebra: On conventional CPU architectures, sparse factorization of these systems dominates the computational cost, severely limiting scalability. These challenges motivate the development of a GPU-based computational framework capable of efficiently handling large-scale nonlinear optimization problems.

3.2 Algebraic Modeling Using ExaModels.jl

The proposed formulation is implemented using ExaModels.jl, which provides a GPU-oriented abstraction for algebraic modeling and parallel automatic differentiation. All nonlinear objective and constraint functions are formulated within this algebraic framework, enabling the automatic generation of first- and second-order derivative information required by Newton-type optimization methods.

ExaModels employs a single-instruction multiple-data (SIMD) parallel abstraction that maps residual evaluation, Jacobian assembly, and Hessian computation directly onto GPU threads. This design enables highly parallel evaluation of derivative information and significantly accelerates the linearization steps performed at each Newton iteration [5].

The hydrogen production and storage models introduced in Section 2 are incorporated into this algebraic framework, enabling a unified representation of the electrical and thermodynamic subsystems within a single optimization model. This unified formulation preserves consistency across coupled physical domains and facilitates efficient execution on GPU architectures.

3.3 GPU-Accelerated Nonlinear Optimization with MadNLP.jl

The resulting nonlinear programming problem is solved using MadNLP.jl, a high-performance primal–dual interior-point solver designed for large-scale and GPU-accelerated nonlinear optimization. At each iteration, MadNLP solves a structured Newton system derived from the KKT optimality conditions.

A central challenge in implementing interior-point methods on GPUs lies in solving the associated KKT linear systems. In standard formulations, the KKT matrix is indefinite and requires sparse symmetric factorization with numerical pivoting to ensure numerical stability. However, pivoting introduces data-dependent branching

and irregular memory access patterns, which are poorly suited for massively parallel GPU architectures. Moreover, as the interior-point iterations approach the solution, barrier terms associated with active constraints lead to increasingly ill-conditioned KKT systems, making purely iterative solvers difficult to deploy robustly.

To address these challenges, MadNLP adopts condensed-space KKT reformulation strategies that transform the original indefinite KKT system into symmetric positive definite (SPD) forms, thereby enabling the use of pivoting-free Cholesky factorizations that are well suited for GPU execution. Two complementary strategies are employed.

MadNLP employs condensed-space KKT reformulations — including Lifted KKT and HyKKT strategies — that transform the indefinite KKT system into symmetric positive definite form, enabling pivoting-free Cholesky factorization suitable for GPU execution [11, 12].

These reformulations allow MadNLP to achieve robust and scalable convergence for highly ill-conditioned multi-period optimization problems with substantial performance improvements over conventional CPU-based implementations.

3.4 GPU-Based Newton Step Computation and Sparse Linear Solvers

The condensed-space KKT reformulations described in Section 3.3 produce SPD systems, which can be efficiently factorized using Cholesky-based sparse direct solvers such as NVIDIA's cuDSS library [13].

Since the condensed KKT matrices are positive definite, cuDSS performs Cholesky factorization without numerical pivoting, avoiding the data-dependent branching that hinders parallel execution. The factorization consists of a symbolic analysis phase — performed only once since the sparsity structure remains fixed across interior-point iterations — and a numerical factorization phase that is reused at each iteration.

The computational workflow proceeds as follows: sparse Jacobian and Hessian matrices are evaluated in parallel using ExaModels and assembled directly into GPU-resident memory, the KKT matrix is factorized using cuDSS, and the Newton direction is recovered via GPU-accelerated triangular solves. Throughout this process, all data remains GPU-resident, minimizing host-device data transfers.

4. CASE STUDY

This section presents numerical experiments to evaluate the computational performance of the GPU-based framework and to analyze the operational characteristics of the integrated power-hydrogen system. The results demonstrate the scalability of the proposed approach and reveal the operational trade-offs arising from

rigorous thermodynamic modeling of hydrogen production and storage.

4.1 Experimental Setup

This subsection outlines the optimization problem formulation, benchmark networks, hydrogen system modeling assumptions, and computational setup adopted in the case study. These components collectively define the experimental environment for evaluating both computational performance and operational behavior.

4.1.1 Optimization Problem

The objective of the case study is to minimize total generation cost over a multi-period planning horizon, subject to electrical and hydrogen operational constraints. The objective function is defined as:

$$\min \sum_{t \in \mathcal{T}} \sum_{i \in \mathcal{N}} (a_i (p_{i,t}^g)^2 + b_i p_{i,t}^g + c_i) \quad (17)$$

where $p_{i,t}^g$ denotes the active power output of generator i , and a_i, b_i, c_i are the corresponding cost coefficients. This objective is minimized subject to the AC power flow constraints (Eqs. 1-5) and the hydrogen system constraints (Eqs. 6-16) detailed in Section 2.

Additionally, a fixed hydrogen demand profile is imposed at each hydrogen bus to represent downstream industrial consumption requirements:

$$h_t^{prod} - h_t^{stc} + h_t^{std} \geq D_t \quad (18)$$

where D_t is the hydrogen demand at time t .

4.1.2 Hydrogen System Specifications

The parameters for the hydrogen subsystems are listed in Table 1. A nominal electrolyzer configuration with a maximum hydrogen production capacity of approximately 1000 kg/h is assumed at each installation site. The storage tank is sized to buffer fluctuations over the planning horizon.

Table 1. Hydrogen production and storage system specifications.

Parameter		Value
Number of cells	N_c	5000
Active cell area	A	2000 cm ²
Hydrogen production capacity	h_{max}^{prod}	1000 kg/h
Operating temperature	T_{elec}	323 K
Hydrogen outlet pressure	P_{H_2}	30 bar
Tank volume	V_{tank}	25 m ³
Operating pressure range	p^{tank}	0-700 bar
Operating temperature range	T^{tank}	270-375 K

4.1.3 Modeling Configurations

Two modeling configurations are considered to assess the impact of thermodynamic detail on both solution accuracy and computational requirements.

The isothermal model assumes constant temperature and pressure throughout the hydrogen process. Adiabatic compression work is approximated using a simplified linear correlation [14], and storage thermal dynamics are neglected. This simplification reduces problem dimensionality and significantly improves computational efficiency.

The thermodynamic model incorporates the full rigorous formulation presented in Section 2. This includes isentropic compression (Eq. 9–10), cooling power requirements (Eq. 11), storage mass balance (Eq. 12), and temperature dynamics via the energy balance (Eq. 14). Tank pressure is determined by the ideal gas law (Eq. 15). This model captures the operational flexibility limits and auxiliary power consumption of the integrated system.

4.1.4 Test Networks and Problem Scale

We evaluate the framework on seven benchmark networks from the PGLib library [15], ranging from the 30-bus IEEE system to a 9591-bus synthetic grid. To simulate a wide deployment of P2H technology, hydrogen production facilities (electrolyzer, compressor, storage) are assumed to be colocated with all active power generators. This placement strategy is motivated by the practical consideration that hydrogen production facilities require access to substantial electrical infrastructure, which is typically available at generation sites.

Table 2. Problem dimensions for benchmark cases with integrated hydrogen systems (60-min resolution). $k=10^3$, $m=10^6$

Case	Periods	Isothermal		Thermodynamic	
		nvars	ncons	nvars	ncons
case30	24	6.0k	8.9k	6.3k	9.1k
case118	24	29.8k	41.8k	32.0k	44.1k
case300	24	68.1k	96.0k	75.0k	102.8k
case793	24	164.0k	227.3k	185.0k	248.3k
case2312	24	490.6k	701.9k	540.2k	751.6k
case4020	24	897.8k	1.4m	908.5k	1.4m
case9591	24	2.0m	3.2m	2.1m	3.2m
case4020	168	6.3m	9.7m	6.4m	9.8m
case9591	168	14.3m	22.2m	14.4m	22.4m

Table 2 summarizes the problem dimensions for each test case under both modeling configurations. We consider two planning horizons: a standard 24-hour day (24 periods) and a week-long 168-hour horizon (168 periods), both with hourly resolution. The inclusion of thermodynamic variables and constraints significantly increases the problem size. For the largest instance (*case9591*, 168h), the thermodynamic formulation involves over 14.4 million variables and 22.4 million constraints, a scale that poses significant challenges for

conventional CPU-based solvers.

4.1.5 Computational Platform

All GPU experiments are conducted on an NVIDIA H200 GPU using MadNLP v0.8.6 as the nonlinear programming solver and cuDSS v0.4.0 for sparse Cholesky factorization. For comparison, CPU-based solvers including MadNLP with UMFPACK and Ipopt with MUMPS are evaluated on a dual Intel Xeon 8480C system (112 cores total, 3.8 GHz max turbo).

4.2 Computational Performance

We evaluate the computational efficiency and scalability of the proposed GPU-native optimization framework (MadNLP-GPU) against two CPU-based alternatives. Figure 2 reports the wall-clock solution times for all benchmark networks under both the isothermal and thermodynamic formulations.

For the isothermal formulation, both CPU-based solvers exhibit reliable convergence on small- to medium-scale instances (up to *case793*). However, their performance degrades rapidly as network size and planning horizon increase. Ipopt successfully solves the 24-hour *case9591* instance in approximately 4,200 seconds, while MadNLP-CPU fails on instances larger than *case793*. In contrast, MadNLP-GPU demonstrates consistently strong performance across all test cases, with solution times ranging from less than 1 second for *case30* to approximately 470 seconds for the largest instance (*case9591*, 168-hour), highlighting the excellent scalability of the GPU-based approach.

The thermodynamic formulation introduces tight bounds on temperature and pressure states, as well as strong nonlinear coupling between electrical and thermodynamic variables. These factors, combined with the wide range of variable scales, significantly degrade the numerical conditioning of the KKT system, consistent with the conditioning challenges discussed in Section 3, and pose severe challenges for CPU-based solvers.

Under this formulation, Ipopt fails to converge for all benchmark instances, and MadNLP-CPU solves only up to *case793*, requiring approximately 3 hours. By contrast, MadNLP-GPU maintains robust convergence across the full test set, solving the largest instance (14.4 million variables, 22.4 million constraints) in approximately 6,650 seconds.

On instances where both GPU and CPU solvers converge, the GPU achieves speedups of 10–600× depending on problem size and model complexity. More importantly, the GPU-based framework extends the tractable problem size: from *case2312* onward, CPU-based solvers fail, whereas MadNLP-GPU scales efficiently to problems with millions of variables. These results confirm that the proposed framework enables high-fidelity modeling of integrated power-hydrogen systems at scales previously inaccessible to conventional approaches.

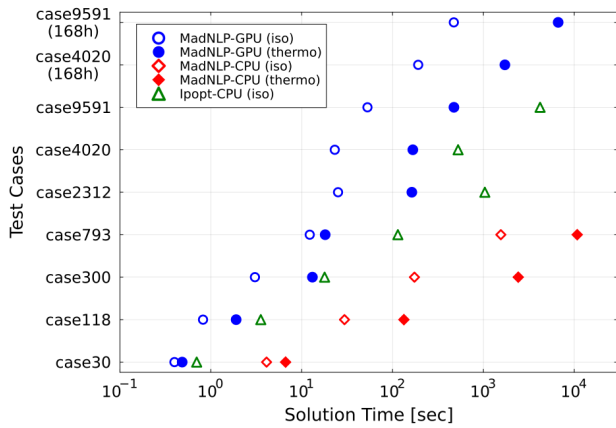


Figure 2. Wall-clock solution times for GPU and CPU solver configurations. Hollow and filled markers denote isothermal and thermodynamic formulations, respectively. GPU refers to MadNLP-GPU. Missing data points indicate solver failure.

4.3 Hydrogen System Operational Analysis

To illustrate the operational behavior captured by the proposed framework, Figure 3 presents the optimal 24-hour dispatch for a hydrogen production facility at bus 69 of the *case118* network. This bus is connected to

a generator and is well-connected to neighboring buses, providing sufficient electrical capacity for hydrogen production. The optimization minimizes total generation cost while satisfying a fixed hydrogen demand of 42 kg/h.

Figure 3(a) shows the hydrogen production rate alongside the baseline electrical load profile. Production varies between 39 and 45 kg/h, exhibiting an inverse correlation with system load. During off-peak hours (hours 1–8), the electrolyzer operates above the required demand at 43–45 kg/h, with the surplus directed to storage. As system load increases (hours 9–10 and 16–23), production decreases to 39–41 kg/h, falling below the fixed demand. This pattern is consistent with both generation cost minimization and network capacity constraints.

Figure 3(b) illustrates the storage dynamics. The tank inventory follows a diurnal cycle, increasing from 50 kg to 60 kg during low-load periods, then decreasing as stored hydrogen compensates for reduced production during high-load periods. The storage thus decouples instantaneous production from the fixed demand.

Figure 3(c) presents the power consumption breakdown. The electrolyzer dominates at 1.86–2.04 MW, exhibiting approximately 9% variation over the horizon. Auxiliary equipment (compressor and intercooler) consumes less than 1 kW under the moderate pressure ratio (30 to 60 bar) considered in this study, representing a

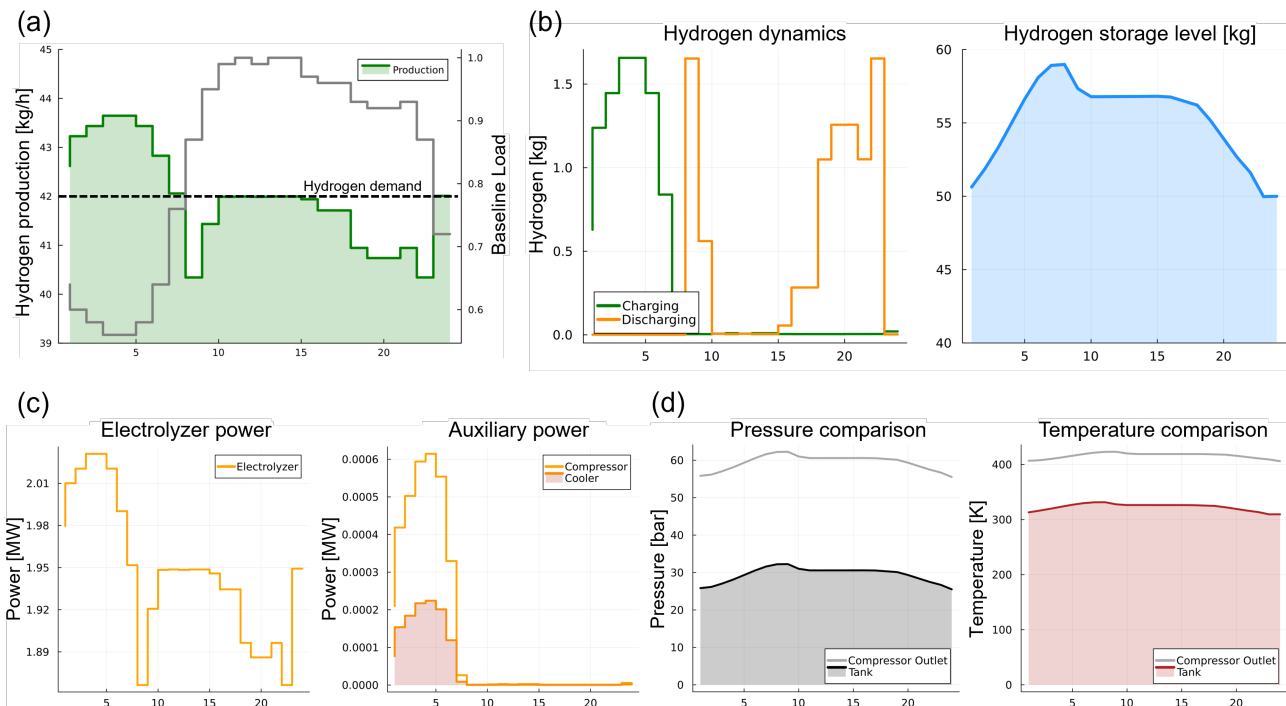


Figure 3. 24-hour operational profiles at bus 69 for the IEEE 118-bus system: (a) hydrogen production rate and electrical baseline load with fixed hydrogen demand, (b) storage charging/ discharging dynamics and inventory level, (c) electrolyzer and auxiliary (compressor and cooler) power consumption, and (d) compressor outlet and storage tank pressure/temperature profiles.

negligible fraction of total consumption. For higher storage pressures (350–700 bar), compression power would become significant, and the detailed thermodynamic model would be essential.

Figure 3(d) shows pressure and temperature profiles. The compressor outlet reaches pressures above 60 bar and temperatures exceeding 400 K during active charging. Following intercooling, the storage tank operates at approximately 30 bar and 300–320 K, within the specified limits (0–700 bar, 270–375 K).

The compressor outlet temperature, exceeding 400 K, must be managed through intercooling to maintain safe storage conditions. Isothermal models would not capture these thermal requirements, potentially yielding schedules that violate equipment limits. This justifies the thermodynamically consistent formulation proposed in this work, demonstrating the practical value of high-fidelity modeling for integrated power–hydrogen system operation.

5. CONCLUSION

This paper proposed and demonstrated a GPU-accelerated framework for multi-period AC-OPF with integrated hydrogen systems, incorporating high-fidelity thermodynamic constraints for electrolyzer, compressor, cooling, and storage operations. The GPU-based solver achieved 10–600× speedups over CPU alternatives, enabling the solution of problems with over 14 million variables, while CPU-based solvers failed to scale beyond medium-sized instances. Operational analysis on the IEEE 118-bus system demonstrated that hydrogen production is inversely correlated with electrical load, with storage buffering production and demand. The thermodynamic model captured critical constraints—such as compressor outlet temperatures exceeding 400 K—that simplified isothermal formulations fail to capture.

Future work will extend this framework to optimal placement of hydrogen facilities, large-scale integration of renewable energy sources, and CO₂ emission mapping for decarbonization-oriented power–hydrogen system planning.

ACKNOWLEDGEMENTS

This research was supported by Korea Institute for Advancement of Technology(KIAT) grant funded by the Ministry of Trade, Industry & Energy(MOTIE), Korea Government (RS-2024-00436106, Human Resource Development Program for Industrial Innovation

REFERENCES

1. Li X, Lepour D, Heymann F, Maréchal F. Electrification and digitalization effects on sectoral energy demand and consumption: a prospective study towards 2050. *Energy* 279:127992 (2023). <https://doi.org/10.1016/j.energy.2023.127992>
2. Shiraishi K, Park WY, Kammen DM. The role of hydrogen as long-duration energy storage and as an international energy carrier for electricity sector decarbonization. *Environ. Res. Lett.* 19:084011 (2024). <https://doi.org/10.1088/1748-9326/ad5856>
3. Singh A, Kumar A, Chinmaya KA, Maulik A. Optimal operation of an electricity-hydrogen DC microgrid with integrated demand response. *Sustainable Energy, Grids and Networks* 39:101451 (2024). <https://doi.org/10.1016/j.segan.2024.101451>
4. Martínez L, Fernández D, Mantz R. Two layer control strategy of an island DC microgrid with hydrogen storage system. *International Journal of Hydrogen Energy* 50:365–378 (2024). <https://doi.org/10.1016/j.ijhydene.2023.09.009>
5. Shin S, Anitescu M, Pacaud F. Accelerating optimal power flow with gpus: SIMD abstraction of nonlinear programs and condensed-space interior-point methods. *Electric Power Systems Research* 236:110651 (2024). <https://doi.org/10.1016/j.epsr.2024.110651>
6. G-L F, Desforges R. Limoges. *essai d'analyse demographique, economique et sociale (etude sommaire)*. *Population (French Edition)* 17:806 (1962). <https://doi.org/10.2307/1526309>
7. Johnson, Sanjay, et al. "ExaModelsPower. jl: A GPU-Compatible Modeling Library for Nonlinear Power System Optimization." *arXiv preprint arXiv:2510.12897* (2025). <https://doi.org/10.48550/arXiv.2510.12897>
8. Baumhof MT, Raheli E, Johnsen AG, Kazempour J. Optimization of hybrid power plants: when is a detailed electrolyzer model necessary?. *2023 IEEE Belgrade PowerTech :1-10* (2023). <https://doi.org/10.1109/powertech55446.2023.10202860>
9. Falcão DS, Pinto AMFR. A review on PEM electrolyzer modelling: guidelines for beginners. *Journal of Cleaner Production* 261:121184 (2020). <https://doi.org/10.1016/j.jclepro.2020.121184>
10. Yigit T, Selamet OF. Mathematical modeling and dynamic simulink simulation of high-pressure PEM electrolyzer system. *International Journal of Hydrogen Energy* 41:13901–13914 (2016). <https://doi.org/10.1016/j.ijhydene.2016.06.022>
11. Pacaud, François, et al. "Condensed-space methods for nonlinear programming on GPUs." *arXiv preprint arXiv:2405.14236* (2024). <https://doi.org/10.48550/arXiv.2405.14236>
12. Regev, Shaked, et al. "A hybrid direct-iterative method for solving kkt linear systems." *arXiv*

preprint arXiv:2110.03636 (2021).

<https://doi.org/10.48550/arXiv.2110.03636>

13. NVIDIA, cuDSS: CUDA Direct Sparse Solver Library, <https://developer.nvidia.com/cudss>
14. Zheng Y, You S, Bindner HW, Münster M. Optimal day-ahead dispatch of an alkaline electrolyser system concerning thermal–electric properties and state-transitional dynamics. *Applied Energy* 307:118091 (2022).
<https://doi.org/10.1016/j.apenergy.2021.118091>
15. Babaeinejadsarookolaei, Sogol, et al. "The power grid library for benchmarking ac optimal power flow algorithms." *arXiv preprint arXiv:1908.02788* (2019). <https://doi.org/10.48550/arXiv.1908.02788>

© 2026 by the authors. Licensed to PSEcommunity.org and PSE Press. This is an open access article under the creative commons CC-BY-SA licensing terms. Credit must be given to creator and adaptations must be shared under the same terms. See <https://creativecommons.org/licenses/by-sa/4.0/>

