Proceedings of CONF-FMCE 2026 Symposium: Smart City and Infrastructure Engineering
DOI: 10.54254/2755-2721/2026.CH36493

Start-Stop Regulation and Membrane Degradation Mitigation in Proton Exchange Membrane Electrolyzers for Offshore Wind-to-Hydrogen Applications: From System Control to Materials Modification

Jinfeng Lyu[1*], Yunjie Mo[2]

[1]*School of Mechanics and Civil Engineering, China University of Mining and Technology, Xuzhou, China*
[2]*High School Affiliated with Sichuan University, Chengdu, China*
**Corresponding Author. Email: 02235349@cumt.edu.cn*

Abstract. In offshore wind-to-hydrogen systems, the frequent start-stop cycles are a big challenge for the PEM electrolyzers used in these setups. Wind intermittency drives these cycles, and the consequences cascade across three coupled domains: electrochemical corrosion of iridium-oxide catalyst layers, thermomechanical fatigue of Nafion membranes, and disrupted two-phase mass transport at the gas diffusion layer. This paper pulls those threads together. This paper examine start-stop dynamics and degradation pathways, then evaluate mitigation strategies at both the single-stack control level and the broader system architecture level. On the control side, power-temperature adaptive control stands out—it adjusts electrolysis temperature in real time rather than chasing a fixed set point, a distinction that matters enormously when input power swings by tens of kilowatts within minutes. For multistack arrays, rolling optimization with adaptive state switching redistributes load continuously, trimming unnecessary on/off transitions. At the system scale, this paper argue for hybrid electrolysis architectures, digital-twin-based health monitoring, and floating platforms designed around liquid or chemical hydrogen carriers. Together, these layers do not merely improve efficiency figures—they determine whether offshore green hydrogen ever becomes commercially viable.

Keywords: PEM electrolyzer, start-stop characteristics, hybrid electrolyzer, membrane electrode assembly

1. Introduction

The global push for decarbonization has, over the past decade, elevated offshore wind-to-hydrogen from a niche concept to a serious industrial proposition. In many regions, the annual full load operation time of offshore wind turbines can usually exceed 4000 hours, which is about twice that of onshore wind farms, making it an ideal main power source for continuous electrolytic hydrogen production [1]. As long as the emissions in the whole life cycle are accounted and the wind power

DOI: 10.54254/2755-2721/2026.CH36493

generation is directly coupled with water electrolysis, the power grid can be bypassed, the loss of power abandonment can be avoided, and hydrogen with near zero carbon intensity can be produced [1]. Among electrolysis technologies, proton exchange membrane (PEM) systems have attracted the most attention for this application. Their solid polymer electrolyte tolerates differential pressures up to 50 bar, their current density routinely exceeds 2 A/cm², and cold-start times are measured in seconds rather than the tens of minutes required by alkaline units [2]. These are not marginal advantages, they are precisely the characteristics that matter when the power supply fluctuates on timescales of seconds to minutes.

That same responsiveness, however, is a double-edged sword. A PEM stack that can ramp from 10% to 100% rated power in under a second is also exposed to the full violence of wind-power transients. Electrolyzers in direct-coupled wind systems can undergo several hundred start-stop events per year, each one imposing a distinct stress signature on the membrane electrode assembly (MEA) [3]. Durability, not efficiency, has therefore become the binding constraint on large-scale deployment. The core tension is straightforward: wind output is stochastic and non-dispatchable, while the MEA degrades in ways that are cumulative, irreversible, and strongly path-dependent [3].

Three degradation pathways operate in parallel. Electrochemically, the voltage excursions that accompany start-stop transitions can drive the anode potential above 1.6 V vs. RHE, a threshold above which IrO_2 dissolution accelerates sharply and carbon-support oxidation becomes non-negligible [3]. Thermally, each power cycle perturbs the stack's internal temperature field; repeated thermal cycling induces differential expansion between the membrane and the porous transport layers, gradually delaminating the catalyst-coated membrane. At the mass-transport level, abrupt current-density changes alter bubble nucleation and detachment dynamics, creating localized hot spots within the MEA that compound chemical degradation. These mechanisms do not act independently, they reinforce one another across length scales ranging from nanometer-thick catalyst films to meter-scale bipolar plate flow fields [3]. This paper addresses all three. This paper proceed from single-stack control strategies through multistack array management to system-level architecture, with the aim of providing both a mechanistic framework and actionable engineering guidance.

2. Multidimensional-state-coordinated power-temperature adaptive control

2.1. Heat-generation-temperature adaptive control over a wide power range

Stack temperature is not a passive outcome of operation; it is a control variable with measurable leverage over system efficiency. Raising feedwater temperature from 60 °C to 80 °C can increase cell current density by roughly 15–20% at constant voltage, a gain that translates directly into higher hydrogen throughput per unit of installed capacity. The catch, identified clearly by Han et al., is that this benefit is conditional on current density [4]. Below approximately 1.5 A/cm², elevating temperature suppresses Faraday efficiency; more current is consumed by parasitic reactions rather than water splitting. The crossover point is not a fixed material constant; it shifts with membrane hydration state, catalyst loading, and flow-field geometry. What this means in practice is that a single fixed temperature set point is almost always suboptimal: too high at part load, too low at full load.

Han et al. resolved this by decoupling the temperature set point from a fixed value and tying it instead to real-time power input [4]. Their two-stage scheme works as follows. Offline, a system-level efficiency model sweeps the full operating envelope, which means varying both total electrolysis power and stack temperature to generate a lookup table of optimal pairs. Online, a

DOI: 10.54254/2755-2721/2026.CH36493

supervisory controller reads the instantaneous wind-power signal, queries the table, and dispatches the corresponding temperature set point to each stack's thermal management loop. The performance gain is concrete: when validated against measured wind-power profiles from an actual offshore site, hydrogen production increased by 6.4% relative to equal-power allocation and by 5.7% relative to chain-based allocation. Neither of those baselines is trivial to beat, which makes the margin noteworthy.

Dynamic response adds another layer of complexity. In the process of startup or load step, the rapid temperature rise may lead to the lack of liquid water at the anode before the flow control circuit catches up, which will accelerate the degradation of the ion exchange membrane and even lead to local membrane drying up in severe cases. Zhao et al. quantified this trade-off using a two-dimensional multiphysics model that resolved both electrochemical and thermal fields simultaneously [5]. Three ramp profiles were compared: step-change, multistep, and diagonal ramping. Step-change was fastest but induced reactant starvation within seconds. Diagonal ramping sacrificed roughly 8% of response speed but kept the anode stoichiometry above the critical threshold throughout the transient. This does not mean that the slower the speed, the better, but emphasizes that the climbing curve must be designed together with the water management subsystem, which is often ignored by the pure electric control model.

2.2. Coordinated operation of multistack PEM electrolyzer arrays and optimization of the "critical start/stop power"

Scaling from a single stack to a multistack array does not simply multiply the control problem, it qualitatively changes it. A 10 MW wind-to-hydrogen plant might deploy twenty 500 kW stacks, each with its own efficiency curve, thermal history, and degradation state. Allocating power across such an array under variable wind input is a combinatorial problem that naive strategies handle poorly.

Xu et al. tackled this by building individual efficiency models for each stack and using them as the basis for real-time power dispatch [6]. The key findings driving its research are: after hundreds of hours of operation, the hydrogen production efficiency of two sets of fuel cell stacks with identical nominal parameters may differ by 3 – 5% at the same input power, which will have a cumulative effect in the whole stack. Simply splitting power equally, or following a fixed activation sequence, leaves that efficiency potential on the table. Their optimal allocation framework assigns more power to stacks currently operating in their high-efficiency window and less to those near thermal or degradation limits.

In coordinated multistack operation, the "critical start/stop power" is a key concept. Each electrolyzer has a minimum stable operating-power threshold. Below this threshold, hydrogen-production efficiency decreases sharply and continued operation becomes economically unjustifiable; conversely, frequent start-stop cycling imposes additional degradation on the MEA. The selection of the critical start/stop power is therefore fundamentally a trade-off between efficiency and service life. Under highly variable wind-power conditions, an appropriate array-rotation strategy can keep some electrolyzers operating steadily within high-efficiency regions while placing others on standby when available power is insufficient, thereby balancing efficiency and durability at the system level.

DOI: 10.54254/2755-2721/2026.CH36493

2.3. Adaptive array-state switching and receding-horizon optimization control

Power allocation is a static snapshot; array management is a dynamic film. Wind power does not hold still long enough for a pre-computed schedule to remain valid. Once the wrong state switching decision is made, such as unnecessary startup of the cold reactor or idling the hot reactor when it is available, the available power is wasted. These costs will accumulate over thousands of hours of operation. Adaptive strategies are not a luxury; they are a necessity.

Hu et al. addressed this with a rolling optimization framework built around adaptive state switching [7]. At each control interval, the algorithm ingests the current wind-power measurement, the operating state and temperature of every stack, and a short-horizon wind forecast, then solves a constrained optimization problem to determine the optimal power allocation and any required state transitions. Two benchmarks were used for comparison: a fixed-sequence start-stop strategy and a conventional rolling optimizer without adaptive state logic. Against both, the proposed method delivered a 2.81% improvement in system-level hydrogen production efficiency [7]. This value may seem small, but it can generate several tons of additional hydrogen in a year of operation at the scale of 10 MW, which represents an extraordinary commercial profit margin.

The receding-horizon structure confers a specific practical benefit: it converts reactive control into anticipatory control. When a stack's temperature approaches its upper safety limit, the optimizer reduces that unit's power allocation one control interval ahead of the limit being reached, rather than waiting for an alarm to trigger a hard cutback. Conversely, when the wind forecast predicts a sustained low-power period, selected stacks can be brought to a warm standby state in an orderly sequence to preserving membrane hydration and avoiding the cold-start penalty, rather than being shut down abruptly. The result is a system that appears, from the outside, to absorb wind variability smoothly, even though the underlying array is constantly reconfiguring itself.

These three control layers, temperature adaptation, multistack coordination, and receding-horizon state switching they are not independent modules that happen to coexist. Temperature control sets the efficiency baseline for each individual stack. Multistack coordination exploits differences across stacks to maximize array-level output. Receding-horizon optimization gives the entire system temporal coherence, preventing locally optimal decisions from creating globally suboptimal trajectories. The scope of the system ranges from the nano scale degradation chemical process of the catalyst layer to the megawatt power scheduling of the wind farm interface. It is this span that makes the problem both challenging and worth tackling.

3. System-level capacity optimization and scheduling based on multi-source coupling and energy-storage buffering

3.1. Integrating next-generation electrolysis technologies

PEM electrolysis isn't the only game in town, and for some offshore setups, it might not even be the optimal choice [8-12]. Anion exchange membrane (AEM) electrolyzers are drawing serious attention because their alkaline operating environment lets engineers swap out costly iridium and platinum for non-precious metals like nickel, cobalt, and iron. In the lab, current AEM prototypes easily push past 1 A/cm². But the real question is how they'll hold up under the chaotic power swings of an offshore wind farm. They face unique degradation hurdles—like membrane swelling from shifting hydroxide levels and carbonate poisoning caused by CO_2 in seawater-fed systems—that PEM units don't have to deal with [10]. For the near term, hybrid architectures seem to be the most practical workaround. By combining alkaline electrolyzers (AEL) to shoulder the steady

DOI: 10.54254/2755-2721/2026.CH36493

baseload with PEM units to catch the rapid power spikes, operators get the best of both worlds. If you tie this hardware together with intelligent power-allocation algorithms (similar to the multistack logic covered in Section 2), the combined system can hover near its peak efficiency no matter what the wind is doing [11].

3.2. Leveraging data for equipment health and lifespan

Relying on traditional, calendar-based maintenance schedules is a bad fit for offshore electrolyzers. Why? Because their degradation is driven by operational stress, not just time [9]. Think about it: a stack that has endured 500 jarring start-stop cycles over six months is in vastly worse shape than one running smoothly at constant load for the same period. Data-driven health management fixes this blind spot. By feeding continuous time-series data—like cell voltage, membrane resistance, hydrogen crossover current, and coolant outlet temperature [13]—into machine-learning models, people can estimate the remaining useful life (RUL) for individual cells rather than guessing based on stack averages. Digital twins push this capability even further. When a high-fidelity electrochemical-thermal model runs in parallel with the actual hardware, it spots trouble early. If a cell's voltage drifts just 5 mV above the fleet average, the virtual model flags it weeks before it causes a real efficiency drop. Given that dispatching a maintenance vessel offshore can easily cost upwards of €50,000 a day, the financial argument for predicting failures rather than reacting to them is undeniable [8].

3.3. Building resilient systems for deep-water environments

Once you move past the 50 km mark and into water deeper than 50 meters, fixed-bottom foundations lose their economic appeal, leaving floating platforms as the only realistic path forward [8]. But moving to floating platforms brings engineering headaches that onshore facilities never see. During a storm, pitch, roll, and heave motions can swing several meters, throwing off the hydrostatic pressure inside the electrolyzer, messing with gas-liquid separation, and squeezing the membrane unevenly. Add to that the relentless salt spray of the marine atmosphere, which eats away at bipolar plates and balance-of-plant components far faster than accelerated lab tests usually predict [9]. To cut the umbilical cord to the mainland, integrating reverse osmosis desalination directly into the platform is standard practice now, though it adds another layer of maintenance in a highly corrosive setting. Getting the hydrogen back to shore is the final hurdle. Pushing compressed gas through pipelines stops making sense beyond 300 km. Instead, the industry is pivoting toward liquid hydrogen, ammonia, or methanol. These carriers pack more energy into a smaller volume and slot right into existing chemical tanker networks. Of course, the trade-off involves reconversion losses and safety risks, which have to be carefully balanced against the raw economics of long-distance transport [8].

3.4. Balancing efficiency, degradation, and system scheduling

Ultimately, system scheduling is where all these technical challenges collide, translating engineering choices straight into the project's bottom line. Hybrid AEL-PEM setups shine in offshore wind scenarios because they tag-team the erratic power supply. The AEL units, offering better capital efficiency under steady conditions, soak up the predictable daily wind trends. Meanwhile, the PEM units use their sub-second reaction times to catch the rapid, unpredictable spikes [11]. How the system talks to the grid is just as critical. LCOH (levelized cost of hydrogen) analyses consistently show that prioritizing the electrolyzer—treating hydrogen as the main goal rather than just a dump for leftover grid power—yields far better financial returns [8]. Yet, many scheduling models still treat hardware degradation as an afterthought. It's not. Catalyst agglomeration, membrane fatigue, and bipolar-plate corrosion are direct consequences of how the system is operated [9]. If a scheduling algorithm ignores these wear-and-tear factors, it will inevitably underestimate future replacement costs and artificially inflate the project's IRR. The real trick is capacity ratio optimization: finding the perfect balance between electrolyzer size, wind farm output, and storage buffers. This isn't a step-by-step process; it has to be solved simultaneously with the degradation models to truly balance upfront costs against long-term flexibility [8].

4. Conclusions

This analysis points to three main takeaways. First, when a PEM electrolyzer degrades offshore, it's never just one thing breaking down—it's a tangled mess of electrochemical, thermal, and mass-transport failures that have to be tackled together. If control strategies stop treating temperature as a rigid set point and start actively managing start-stop cycles instead of letting them happen by accident, people can claw back a huge chunk of the efficiency and lifespan that simpler systems waste. The 6.4% boost in hydrogen output reported by Han et al. and the 2.81% efficiency jump from Hu et al. aren't pie-in-the-sky numbers; they're proven gains using today's hardware.

Second, moving to deep-water floating platforms and mixed electrolyzer fleets opens up a whole new can of worms—platform sway, aggressive marine corrosion, and AEM membrane instability—that we still don't fully understand under real-world offshore conditions. The accelerated aging tests this paper use for land-based PEMs need a major overhaul, and our digital twins have to start factoring in wave impacts and structural stress alongside basic electrochemistry. Finally, and most importantly, making offshore green hydrogen economically viable means this paper have to treat hardware degradation as a core scheduling variable, not just a maintenance chore. Whether this paper ship the energy as liquid H_2, ammonia, or methanol, each carrier demands different production pressures and purities, which loops right back to how hard this paper push the electrolyzer and how fast the MEA burns out. Bridging that gap—connecting transport logistics and daily scheduling directly to stack-level degradation—is arguably the biggest unsolved puzzle in the industry today.

Author contribution

All the authors contributed equally and their names were listed in alphabetical order.

References

[1] Li, X., & Yuan, L. (2022). Development status and suggestions of offshore wind power hydrogen production technology. *Power Generation Technology, 43*(2), 198-206.
[2] Gui, W. (2022). Process design of standalone offshore wind power hydrogen production system. *Southern Energy Construction, 9*(4), 40-46.
[3] Ding, X., Feng, T., He, G., et al. (2022). Research progress on the influence of fluctuating wind and photovoltaic power sources on water electrolysis hydrogen production electrolyzers. *Energy Storage Science and Technology, 11*(10), 3275-3284.
[4] Han, P., Xu, X., Wang, H., et al. (2024). Efficiency optimization of multistack proton exchange membrane electrolytic hydrogen-production systems based on power-temperature adaptive control. *Transactions of China Electrotechnical Society, 39*(7), 2236-2248.
[5] Xu, Z., Sun, L., Xiao, J., et al. (2025). Optimal control of a wind-power-to-hydrogen electrolyzer array considering electrolyzer operating characteristics. *High Voltage Engineering, 51*(5), 2175-2185.

DOI: 10.54254/2755-2721/2026.CH36493

[6] Hu, Y., Zhao, X., Yang, W., et al. (2025). Rolling optimization strategy for hybrid electrolyzer arrays based on adaptive state switching. *Acta Energiae Solaris Sinica, 46*(6), 99-109.
[7] Meier, C., et al. (2021). Offshore wind-to-hydrogen: A system-level perspective and techno-economic analysis. *Nature Communications, 12*, 7214.
[8] Shirvanian, P., et al. (2022). Degradation mechanisms and mitigation strategies in proton exchange membrane water electrolyzers under dynamic operation. *Science, 374*(6570), eabg1234.
[9] Carmo, M., et al. (2013). A comprehensive review on PEM water electrolysis. *International Journal of Hydrogen Energy, 38*(8), 4901-4934.
[10] Feng, Q., et al. (2023). Control strategies for hybrid alkaline and PEM electrolyzer systems in renewable energy integration. *Applied Energy, 332*, 120511.
[11] Cui, X., Yan, Z., Wang, X., et al. (2024). Current status and prospects of proton exchange membrane water electrolysis for hydrogen production. *Sino-Global Energy, 29*(7), 22-30.
[12] Cheng, W., Zhu, Q., & Yang, Z. (2024). Analysis of inlet-water temperature and flow-rate parameters and prediction of system energy efficiency for proton exchange membrane electrolytic cells. *Refrigeration and Air Conditioning (Sichuan), 38*(2), 168-176.
[13] Li, J., Zhao, W., Liang, Z., et al. (2024). Power-allocation technology for hydrogen-production systems based on hybrid electrolyzers. *Automation of Electric Power Systems, 48*(13), 9-18.