H2UpClose #8: RWTH Aachen University

Inside the Minds Driving Hydrogen Innovation

Version: 1.0
YYYY–MM–DD
HORIZON-JTI-CLEANH2-2024-03-02

H2UpScale

Clean Hydrogen Partnership

Co-funded by the European Union

H2UpClose #8: RWTH Aachen University

H2UpScale Project is happy to present the H2UpClose interviews' cycle.
Below the eighth interview with RWTH Aachen University.

In one sentence, what problem are you obsessed with solving in this project?

How to design and optimize a fuel cell system at a power level where the supporting components, the compressor, the recirculation blower, the cooling, all the balance-of-plant, stop being accessories and start defining whether the whole system works, and then how to operate several of these high-power systems together with a battery so the powertrain is efficient at every moment while protecting the systems from unnecessary ageing.

Finish this: "If we get this thing right, ______ will change for ______."

If we get this right, two things change for the engineers and operators scaling fuel cells to maritime power levels. Balance-of-plant sizing moves from something you discover by building and measuring hardware to something you decide in simulation before you commit. And energy management in multi-fuel-cell powertrains moves from simply meeting the requested power to actively deciding which systems run, how much each provides, and how the battery supports them, in service of efficiency, lifetime, and predictable operating cost.

What's the most surprising thing people get wrong about Fuel Cells' upscaling?

That you can scale by scaling the stack. The intuitive picture is that a bigger fuel cell is mostly a bigger stack, with the balance-of-plant sized up proportionally. It's the opposite: the stack is in some ways the well-understood part, and it's the balance-of-plant that doesn't scale cleanly. Parasitics grow faster than the power, thermal management gets disproportionately harder, and the components you'd want often don't exist at these sizes yet. A second misconception follows once several systems are installed: that they should share the load equally. They shouldn't. Individual systems age differently and have different efficient operating points, so our energy management distributes load by their condition and ageing state rather than treating them as identical.

What's the "before vs after" story?

The conventional way to size these systems is to bolt existing components together, build a prototype, measure it, and iterate in hardware. That's slow and expensive, and at the 350 kW to 1050 kW levels we're targeting the components you'd want often don't exist off the shelf yet, so you're partly guessing. We instead build a physics-based dynamic model of the whole system with an optimization framework on top, so sizing decisions get made in simulation against real use-case demands before anyone commits to hardware. The same shift applies to operation, where energy management is mostly rule-based today: we develop a two-layer EMS, a supervisory layer deciding which systems are active and a model predictive control layer optimizing the power split, so the strategy delivers the required power at higher efficiency and reduced degradation.

Why does solving this problem matter now and not in five years?

Because the components and the architecture are being decided right now. Heavy-duty fuel cell systems are moving from the few-hundred-kilowatt research scale toward commercial and maritime power, and the choices locked in today, how big the compressor is, how thermal management is laid out, how anode recirculation is done, will shape cost and efficiency for the next decade. The same is true one level up: once several systems are combined with a battery, the question shifts from how to design the components to how to operate them over thousands of hours, where the operating strategy directly drives hydrogen consumption, degradation, and total cost of ownership. Develop good sizing methods and ageing-aware EMS strategies now and the field scales faster; leave it to hardware trial and error and we lock future systems into inefficient, lifetime-reducing concepts.

What are the top two constraints that make this hard?

First, the balance-of-plant doesn't scale linearly. As you push to higher power the air supply dominates the parasitic load: scaling mass flow and pressure ratio pushes the turbomachinery into regions where off-the-shelf components lose efficiency, so parasitics grow faster than the power. Thermal management also gets difficult, because a PEMFC runs at around 80°C with no hot exhaust to carry heat away, so almost all the waste heat goes through the coolant loop and out a radiator at a small temperature difference to ambient. Second, the right sizing depends entirely on the duty cycle, which sets the load dynamics and time spent at each operating point, so the optimum is tied to the use case rather than any single "best" design. The EMS faces a parallel pair: short-term efficiency is coupled to long-term ageing, and it must coordinate the battery and several systems at once across state of charge, operating limits, transient demand, start-stop decisions, and individual ageing. That makes it fundamentally different from single-stack control: we're deciding how a fleet of systems inside one powertrain should cooperate. applications.

What happens if industry does nothing about it?
You get systems that are oversized, inefficient, or expensive, or all three, because the safe choice without good models is to over-provision, which makes fuel cells less competitive than they need to be. The operational side mirrors this: without an ageing-aware EMS, modular powertrains get run in ways that look acceptable short-term but are costly over the lifetime, with some systems ageing faster than others, higher fuel consumption, and earlier component replacement, and so more material demand as drivetrains are replaced sooner.

What is your role and what are your goals in the project?
We're RWTH Aachen University, and our role spans two layers. The first is system modeling and balance-of-plant sizing optimization: we build a full dynamic simulation of the fuel cell system, every major component in both the air and hydrogen loops plus thermal management and the dynamics of how they interact, then wrap a multi-criteria optimization framework and cost models around it, so we can ask what the right size is for each component given a target power and use case. The second is the energy management system for the hybrid multi-fuel-cell powertrain: an EMS that distributes requested power between the battery and the individual systems based not only on instantaneous efficiency but on each system's ageing state.

What piece of the puzzle is your team uniquely positioned to deliver, and why you?
On sizing, we optimize three criteria simultaneously: efficiency, performance over a use-case cycle, and cost. They genuinely pull against each other, a bigger compressor improves performance and lets the stack run more efficiently but costs more and draws more parasitic power. Rather than collapse this to a single number too early, we trace out the trade-off front so a decision-maker sees what a given cost buys and chooses deliberately. The EMS has its own central trade-off between instantaneous efficiency, battery usage, and long-term degradation: running a system near its efficiency optimum helps hydrogen consumption but ages it faster if overused, while sharing power equally balances usage but moves systems off their best efficiency points.

What added value does your activity bring?
On sizing, we optimize three criteria simultaneously: efficiency, performance over a use-case cycle, and cost. They genuinely pull against each other, a bigger compressor improves performance and lets the stack run more efficiently but costs more and draws more parasitic power; cheaper thermal management saves cost but compromises efficiency. Rather than collapse this to a single number too early, we trace out the trade-off front so a decision-maker sees what a given cost buys and chooses deliberately. The EMS has its own central trade-off, between instantaneous efficiency, battery usage, and long-term degradation: running one system near its efficiency optimum helps hydrogen consumption but ages it faster if overused, while sharing power equally balances usage but moves systems off their best efficiency points.

What trade-off did you optimize (performance, cost, reliability, safety) and how?
You get systems that are oversized, inefficient, or expensive, or all three, because the safe choice without good models is to over-provision, which makes fuel cells less competitive than they need to be. The operational side mirrors this: without an ageing-aware EMS, modular powertrains get run in ways that look acceptable short-term but are costly over the lifetime, with some systems ageing faster than others, higher fuel consumption, and earlier component replacement, and so more material demand as drivetrains are replaced sooner.

Which society or market long-term goals will H2UpScale help achieving?
Decarbonizing the parts of transport hardest to electrify with batteries alone. I'll be straightforward: for many heavy-duty road applications battery-electric vehicles are likely to dominate, and that's a good thing. But there are use cases, certain demanding long-haul or high-utilization road profiles, aviation over a longer horizon, and especially shipping, where the energy density and fast refueling of hydrogen and fuel cells make them a well-suited technical solution. To serve those, fuel cell systems have to scale to much higher power than today's automotive systems, and getting the system design and balance-of-plant right is what makes that scaling efficient and affordable enough to deploy.

Is the cooperation with international Partners in a unique framework an added value in your opinion? And why?
Very much so. Scaling a fuel cell system isn't one discipline's problem; it touches stack technology, turbomachinery, filtration, power electronics, vehicle and vessel integration, and the modeling that ties it together, and no single organization holds all of that. In H2UpScale you have OEMs across different applications, component specialists, and research institutions in one consortium, so the system-level work we do is grounded in real component data and integration constraints rather than assumptions. For modeling specifically, that connection to people who actually build the hardware is what keeps the models honest.

Where can people follow progress or get involved right now?
People can follow progress and get involved by visiting the H2 Upscale project website, connecting with consortium partners on LinkedIn, or subscribing to project updates and newsletters shared by the participating organizations.

Website: www.h2upscale.eu
LinkedIn: https://www.linkedin.com/company/h2upscale

Acknowledgement
The project is supported by the Clean Hydrogen Partnership and its members.