Green Hydrogen: Some Key Numbers

PART I: Energy Efficiency

Introduction: Why do we need clarity about the limits of Hydrogen?

Hydrogen is perhaps among the hottest investment topics in Climate Tech but it is also probably the most controversial one. There has been a continued raging debate whether we should develop a hydrogen economy using hydrogen as a Swiss Army Knife to solve almost all kinds of decarbonization needs. Spurred by a period of virtually free money, government incentives and support, commitments to net zero, there was a frenzy rush in scaling up hydrogen with massive commitments by both governments and private sectors. By 2023, the combined projected capacity of the top 20 Hydrogen Projects added up to a production of 32 million tonnes/year (for reference, the world is currently producing about 100 Mt H_2/year). Only ~4% of all *announced* electrolyser capacity for 2030 had reached FID (Financial Investment Decision) or construction as of the IEA's 2024 review (even then FID has a new meaning as it does not necessarily signal 100% commitment to move forward). Over the last several months, there has been a devastating flood of news about hydrogen project cancelations. For a non-exhaustive list of canceled hydrogen projects, see here: **Cancelled and postponed green hydrogen projects | Reuters**

Hydrogen bubble deflated as fast as it went up. WHY? This situation raises the question: how can such enthusiasm shift to disillusionment so quickly? When market fundamentals are robust, dramatic reversals are unusual. For the non-specialist, it can be challenging to comprehend how something that was so certain of a major role and got so much commitment and support just suddenly literally got the rug pulled under it. Most importantly, that rush of capital to support hydrogen development followed by the rush to cancel projects, often with billions of dollars in subsidies, is so telling of a panic wake up, so wasteful (we should also mention the impacts they have on people life such as the 700 jobs losses in the Fortescue H_2 project) and so damaging to the general public trust that lessons should be drawn. The prevailing explanation for this abrupt shift centers around cost and infrastructure leading to insufficient offtakes to justify the projects. However, it is concerning that these challenges were only recognized post-investment. What was missing? This type of mistake is unfortunately way too common especially in the startup world, "just build it and they will come", but the offtakes did not come because the price of hydrogen even with subsidies was unaffordable. The wild dream just got confronted with harsh reality. From a technical standpoint, longstanding concerns have been raised about hydrogen's viability due to thermodynamic limitations making it unsuitable for many applications, especially as a fuel. While such technical reasoning may be difficult for non-experts to grasp, it remains central to understanding the sector's challenges.

A lack of understanding of the true reason(s) could mean that investment decisions were made on the wrong basis, causing the herd behavior in the 2020s, AND subsequent project cancellations were also made on the wrong basis when the herd ran for an exit. This pattern raises concerns that such errors could persist and potentially recur when investment conditions become favorable once more. Proof point: many experts, including Ulf Bossel during Cleantech 1.0, warned back then that hydrogen is an inefficient energy carrier, with electricity from hydrogen and fuel cells costing at least four times more than grid electricity. Two decades later, the industry still lacks adequate responses to Bossel's well-founded critique. Nevertheless, a renewed push during Cleantech 2.0 saw even larger-scale investments, only to be curtailed very rapidly by economic realities: high costs deter buyers. Absent a thorough understanding of the fundamental causes behind these setbacks, there is still a sense of lingering hope that, given enough investment, time and scale, hydrogen cost can come down. This argument is often bolstered by the precedent set by solar and battery technologies, which experienced spectacular cost declines over time. Hope is powerful to make big things happen, but false hopes can also be draining and could become an expensive opportunity cost, in the sense that the investment and time spent on developing the hydrogen economy could be spent on building something more effective, the electron economy for example. Already, the slowdown with hydrogen development has been interpreted by a few as a healthy shakeup to eliminate projects that don't have strong fundamentals. Maybe Fortescue, BP, Shell, Iberdrola, Air Products, Trafigura, ArcelorMittal, Cleveland Cliffs, etc. did not get it right? Accordingly, a critical question moving forward is whether green hydrogen merits another opportunity to achieve economies of scale and cost reductions. Unless there is a clear scientific and economic assessment of hydrogen's ultimate potential for performance and cost improvements, hope will continue and may lead to further investments with unproven returns.

In Part 1, we conduct an in-depth energy analysis to clarify the limitations of hydrogen, focusing on whether the primary constraint is thermodynamics and examining if improvements similar to those made with solar are possible. Part II, to be published subsequently, will evaluate the implications and impacts of using green hydrogen as an energy carrier for power generation or heating fuel, considering cost, CO_2 emissions, effects, and any unintended consequences to address the essential question that all new technologies aspiring to solve problem should answer: is the cure better than the disease? Part I contains technical scientific content presented at a high level to explain the physical limitations. The aim is to offer clear, simplified, science-based explanations in accessible terms so that individuals from various backgrounds can make informed decisions.

Summary for Part I: Thermodynamic limitations are not the primary factor leading to electricity-hydrogen-electricity conversion efficiency below 30% as typically described in the literature; rather, low efficiency is primarily dictated by CAPEX requirements. Nonetheless, thermodynamics ultimately constrain hydrogen's potential due to the fundamentals of H_2/H_2O redox electrochemistry, which cannot be altered by additional research and development. Except for certain specific applications, green hydrogen represents a wasteful use of renewable energy when used as a fuel.

Basic understanding of hydrogen energy

Interestingly, hydrogen reaction with oxygen is so simple on paper, but its energetic balance is not, and this could be one of the reasons for confusion and misunderstanding.

The electrochemical reaction of hydrogen in a fuel cell can be written as:

$$H_2 + 1/2O_2 \leftrightarrow H_2O_{(liquid)} + 237.2\text{ kJ/mole } \textit{electricity} + 48.6\text{ kJ/mole } \textit{heat}$$

237.2 kJ/mole is the Gibbs free energy for the reaction and therefore the maximum electricity that can be produced by the fuel cell. The reaction also generates heat, which corresponds to the energy loss due to entropy change (TΔS in thermodynamics), meaning that this heat energy cannot be converted to useful work. Heat could be recovered as low-grade waste heat by condensing steam back to water but that is limited to specific applications. The sum of both electrical energy and heat energy corresponds to the enthalpy of reaction ΔH, which is also called the Higher Heating Value (HHV) of hydrogen. The Gibbs free energy corresponds to a thermodynamic equilibrium voltage of 1.23 V while that of the enthalpy-HHV is 1.48 V. The correct unit of energy should be in electronvolt (eV) rather than just volt (V), but for simplicity of discussion, we will use volt unit with the understanding that eV should be kept in mind for charge conservation when connecting an electrolyzer with a fuel cell.

The electrolysis reaction is exactly the opposite of the fuel cell reaction (equation (1) going from right to left). This means that both electrical energy and heat must be provided to drive the reaction toward hydrogen generation from water. This means that in electrolysis mode, heat from an external source could be added to the reaction in addition to electricity to split water. However, heat is typically generated in-situ since the electrolyser tends to have sufficient internal resistance to cause far more heat generation than what is required.

The definition of energy efficiency is defined as the ratio of what is going out that is useful over the input needed to make things happen, which means (Electricity Out)/(H_2 Energy In) in fuel cell mode and (H_2 Out in energy content)/(Energy In) in electrolysis mode. The theoretical maximum electrical conversion efficiency for a fuel cell is 1.23/1.48 = 83%. For the electrolyzer, the theoretical energy efficiency of an electrolyzer is 1.48/1.48 = 100%, while the electrical conversion efficiency can be up to 1.48/1.23 = 120%, simply because in principle, external heat can be brought in to help drive the electrolysis; but this has no practical application. As a result, the Round-Trip-Efficiency (RTE) of using electricity to make hydrogen by electrolysis, then using hydrogen as energy carrier to make electricity again is the product of the two HHV efficiency numbers, i.e. 83%, a quite decent number in theory. Up until this point, thermodynamics are responsible for the 17% drop in efficiency.

Basic electrochemistry: The reality of electrochemical energy conversion

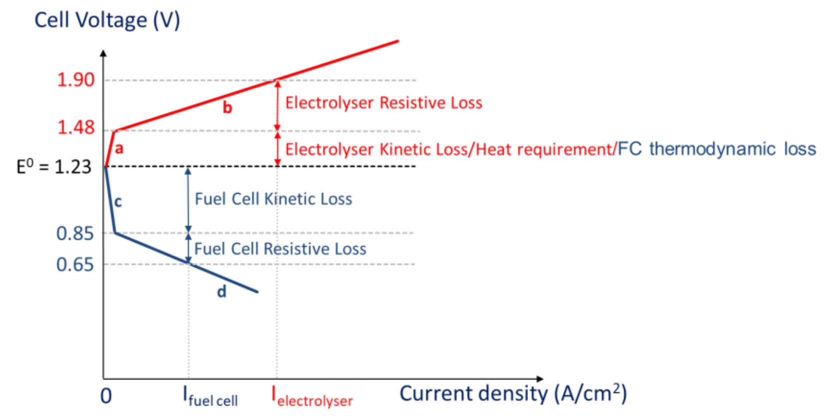

FIG. 1: Schematic of Current-Voltage curves for PEM electrolyser and fuel cell

Practically, electrolysers and fuel cells never operate at or even near the theoretical efficiency points because of real-world losses. This is where things unravel quickly. The schematics of the current-voltage curves for both the electrolyser (in red) and the fuel cell (in blue) are shown together in FIG.1 to assess the round-trip energy transformation from electricity to hydrogen and back to electricity. We will use the example of Proton Exchange Membrane (PEM) electrolyser and PEM fuel cells for discussion purposes because they are the most advanced technologies to date. However, similar conclusions can be drawn for other low temperature technologies including alkaline electrolyzers.

As seen in FIG.1, the electrolyser and the fuel cell can operate over a wide current density range. Let's consider the electrolyser first: the thermodynamic equilibrium voltage is 1.23 V (E^0), which implies that a voltage higher than 1.23 V is needed to drive the water electrolysis reaction. As seen in FIG.1, the voltage increases rapidly for a very small increase in current density (zone a) resulting in a steep initial slope. This is because substantial energy is initially needed to overcome the reaction barriers at the electrodes, something that is called kinetic loss. The kinetic loss generates heat that otherwise would need to be provided from an external heat source to support the electrolysis reaction. Above 1.48 V, the curve changes to a much smaller slope, zone b, where various cell components in the electrolyser cell (membrane, contacts, etc.) are impeding the current flow sufficiently in form of resistance, causing resistive losses. For the fuel cell, again, the thermodynamic equilibrium voltage for electrical generation is 1.23 V. As soon as current is drawn, voltage also drops rapidly in zone c between 1.23 V all the way down to about 0.85 V for similar reasons as for the electrolyzer, i.e. due to kinetic loss. When the kinetic barrier is sufficiently overcome, cell resistance starts to dominate the voltage loss at voltages below 0.85 V and the curve slope is much less steep (zone d). A note for the specialists in electrochemistry, there is a contribution from mass transport to voltage losses in zone d because fuel cells must operate close to 80% fuel utilization for efficiency reasons, but I will not discuss it here for simplicity purpose.

In theory, the electrolyser and the fuel cell can operate at any current density greater than zero, i.e. at any point on the x axis of FIG.1. However, manufacturers typically design them to operate around a design current density that is far away from the zone dominated by kinetic losses. The design current density is a compromise between two opposite requirements. On one hand, efficiency should be kept as high as possible, i.e. the cell must operate at as low a voltage as possible for the electrolyser, or as high a voltage as possible for the fuel cell, which implies keeping current densities as low as possible for both the electrolyser and fuel cell. On the other hand, CAPEX is another major consideration, especially because there are expensive components in the cells including precious metals such as platinum and iridium. Operating at too small current densities results in unacceptably large size cells and unaffordable high CAPEX. In other words, electrolysers and fuel cells must operate at as high current density as possible to minimize CAPEX. This is the reason why electrolysers and fuel cells operate in the A/cm^2 range vs mA/cm^2 for batteries. Typical voltage for a commercial PEM electrolyser is between 1.8 to 2 V, a far distance from the thermoneutral voltage of 1.48 V; for the discussion here, we will pick a representative number of 1.9 V. Typical operating voltage for a commercial PEM fuel cell is between 0.7 to 0.6 V; we will pick a representative number of 0.65 V, which is the most common number. This is also a far distance from the thermodynamic equilibrium voltage of 1.23 V. Note that PEM electrolyzers and PEM fuel cells don't necessarily operate at the same current density, with the electrolyzer typical current density being around 2 to 4 A/cm^2, while that of a PEM fuel cell being closer to 1 A/cm^2.

Real world Efficiency

With real world values for fuel cell and electrolyser operating voltages, now we can calculate real world efficiencies. As mentioned above, efficiency is defined as what we get out divided by what we put in; in this case, it is as simple as fuel cell voltage/electrolyser voltage (should be in electronvolt), i.e. 0.65/1.9 = 34%. This efficiency can also be described as the product of the 2 steps:

$$RTE = Electrolyer\ Eff * Fuel\ Cell\ Eff = \frac{1.48}{1.9} * \frac{0.65}{1.48} = \frac{0.65}{1.9} = 34\%$$

So far, this RTE number only concerns the DC voltages of the electrolyser and the fuel cell, we have not yet included all system components & intermediate steps. Yet, it is already a very low number, far below the theoretical efficiency of 83%. What causes that massive drop in efficiency? The efficiency loss for each step in the electrochemical conversion can be calculated from the voltages found in FIG.1. Table 1 shows the percentage contributions of various factors to the total efficiency loss. Interestingly, thermodynamics contribute only 20% to the total loss. The desire to operate the electrolyser and fuel cell at practical current densities to minimize footprint and CAPEX is causing the remaining 80% loss.

Break down of voltage loss	**V loss (V)**	**% of total loss**
Elyzer kinetic loss/Elyzer heat/FC tΔS	0.25	20%
Elyser resistive loss	0.42	34%
FC kinetic loss	0.38	30%
FC resistive loss	0.2	16%
Total	1.25	100%

Table 1: Breakdown of various voltage loss contribution

To enable the conversion of electricity to hydrogen and then back to electricity, many more components and steps must be included in the round-trip transformation as pointed out in numerous references in the literature. The list includes the AC-DC converter and Balance-of-Plant (BOP including pumps, controls, etc.) for the electrolyser, hydrogen compression system to compress H_2 from 30 bars (typical pressure coming out of PEM electrolysers) to 700 bars for storage, optional shipping and/or delivery system of that hydrogen to the point of use, AC-DC converter and BOP for the fuel cell. In some applications where higher volumetric energy density is required for storage such as transoceanic shipping, etc., hydrogen is liquefied instead of being compressed to 700 bars. The transformation of that liquid hydrogen back to a gas before use may require additional heat (i.e. energy) input that is not included in the analysis. It has been highlighted before that making a very diffuse gas like hydrogen into a form practical to use, i.e. dense enough, requires a lot of energy, about 3 – 4 kWh/kg H_2 using compression and as much as 10 kWh/kg H_2 or 30% of the energy content in hydrogen using liquefaction (not counting continuous losses due to regular venting). As shown in FIG.2, when all those additional steps are included beyond the electrolyser and fuel cell stacks, the RTE drops from 34% to 30% for compressed H_2 and to 26% for liquefied H_2 (not shown). Note that there are more losses that are not even counted in the analysis for simplicity purpose, including for example that hydrogen fuel utilization in a fuel cell is not 100% even with recirculation, which results in some efficiency penalty. What is most interesting is, the new numbers don't appear as bad as they should be based on concerns about handling of hydrogen gas. This is mainly because the starting point, 34% (not 100%), is already so low that additional losses as a percentage of the starting point don't turn the number terribly lower. As seen in FIG.2, the thermodynamic losses, which include hydrogen compression and $T\Delta S$ heat (shown in red), are only responsible for a combined 19% drop in efficiency, while other non-thermodynamic factors are responsible for a 51% drop.

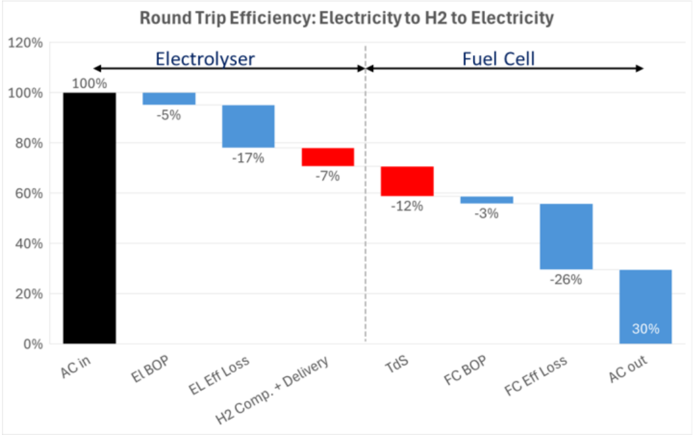

FIG. 2: H2 RTE when counting supporting system components

Can hydrogen efficiency be improved with more development and time?

Certainly; but a more pertinent question is how much and would that matter? Since the PEM fuel cell has been in development for 60 years since NASA's Gemini program, meaningful further improvement in performance is unlikely, so we will not adjust fuel cell operating voltage. FIG. 3 shows the impact of electrolyser operating voltage on RTE. Even if voltage could be lowered all the way down to the thermoneutral voltage (note that there is a claim out there of a breakthrough 95% electrolyser efficiency, i.e. operating at 1.56 V; I will believe it when we get to see a commercial stack, not a lab cell, capable of operating at that level) which is the limit beyond which external heat must be added to support the reaction, RTE is still below 38% for compressed hydrogen and 33% for liquefied hydrogen in the best-case scenarios. This is the ultimate limit of using hydrogen as an energy carrier to convert back to electricity. Whether this efficiency is good enough can be debated; however, in absolute terms, these values are still too low if we define anything below 50% as not good enough (an arbitrary feel-good threshold inspired by lithium-ion battery competition which is at 90%), resulting in too much waste of precious renewable electricity. Therefore, just from efficiency consideration, the potential gain/benefit from getting closer to the fundamental limit of electrolyser performance is so small that it may still not justify the push for more hydrogen development and deployment. Consequently, for electricity generation using electrolytic hydrogen fuel, higher efficiency does not necessarily win, because the win is still not enough to make a meaningful change in overall system efficiency. Note that lowering electrolyser voltage will likely come at higher CAPEX cost because there is no free lunch, so the net win may be even smaller.

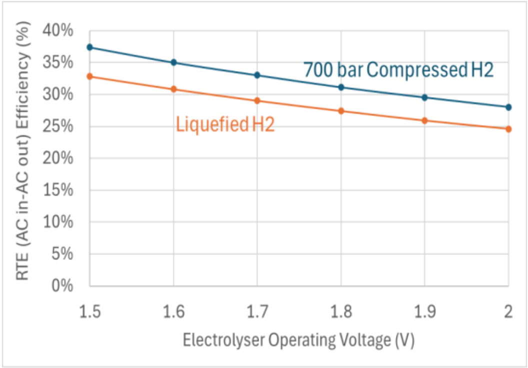

FIG. 3: What is the opportunity for improvement?

How about using a combined-cycle gas turbine instead of a fuel cell for power generation as often proposed for long-duration energy storage schemes? The efficiency of a combined-cycle gas turbine can be better than a fuel cell, around 50% HHV instead of 44%, this would bring the RTE from 30% to 34% for compressed hydrogen (but that comes with new technical challenges including high NOx emissions). Does that really change the picture in a meaningful way?

How did we even get to accept such low efficiency to start with?

30% RTE may sound like a bad number, but it was not bad not too long ago. This is because we should always compare anything new to a reference to beat. Years ago, during Cleantech 1.0 wave, hydrogen was compared to the internal combustion engine (ICE) in vehicles. The efficiency of ICE using gasoline is about 20%. Hence, 30% was a convincingly better number in addition to the main benefit of having zero emissions.

But time has changed since and the goal post has moved. Batteries are now the reference to beat, whether they are used in Electric Vehicles (EVs) or for energy storage. What is battery AC to AC RTE? about 90%, with the added benefit of a much simpler system (which will result in lower total system cost). At 30%, hydrogen RTE is thus 3x lower than the dominant competitor. Again, this is not due to thermodynamics, at least not in the traditional sense of the complex multi-step conversion.

But low efficiency is due to fundamental thermodynamics when compared to lithium-ion batteries!

The process of converting electricity to hydrogen and back to electricity shares similarities with the charging and discharging cycles in batteries. In this context, the electrolyser functions during the “charge cycle” by producing hydrogen for storage, while the fuel cell operates during the “discharge cycle” by consuming hydrogen to regenerate electricity. Unlike batteries, which exhibit variable voltage relative to State-of-Charge (SOC) throughout these cycles, both the electrolyser and fuel cell maintain constant charge and discharge voltages due to continuous water or hydrogen supply.

Both hydrogen systems and batteries operate near their thermodynamic equilibrium voltages, that are 1.23 V for hydrogen, and 3.7 V for lithium-ion batteries at 50% SOC. As previously noted, energy efficiency should be understood as a ratio rather than an absolute value. The principal distinction between lithium-ion batteries and hydrogen lies in the denominator of this efficiency calculation: the available energy to perform work, which is three times greater for batteries. This elevated voltage characteristic derives from lithium chemistry and is facilitated by the use of organic electrolytes that allow for a wider operating voltage range, whereas the hydrogen-water system is restricted by the decomposition threshold of water at 1.23 V.

Consequently, independent of conversion losses, lithium-ion batteries possess an inherent efficiency advantage over H_2/H_2O systems based purely on chemical thermodynamics; any loss, when divided by a denominator three times larger, results in a proportionally smaller percentage loss. *This fundamental property cannot be changed through further research and development.* Additional real-world inefficiencies associated with electrolyser and fuel cell operations—typically higher than those observed in batteries due to more complex conversion pathways—further reduce overall efficiency but do not alter the underlying thermodynamic limitation of hydrogen compared to lithium-ion technologies.

Is there a better use of hydrogen?

It should be obvious by now that the use of hydrogen as a fuel—whether for fuel cells or gas turbines—represents the primary source of energy inefficiency in these systems. As illustrated in FIG. 2, the hydrogen production process achieves an efficiency of 74% at the output of the electrolyser, or more realistically, 68% when accounting for the compression required for hydrogen storage. It is important to note that storage is invariably necessary, as synchronizing the mass flow requirements of downstream applications with the output from the electrolyser is seldom feasible. By utilizing hydrogen solely as a chemical, rather than using it as a fuel for energy production, losses associated with entropy ($T\Delta S$) and the reconversion of hydrogen to electricity can be avoided. This rationale underpins the concept of the hydrogen application ladder as proposed by Michael Liebreich ((2) Hydrogen Ladder Version 5.0 | LinkedIn), which suggests that employing hydrogen in chemical processes—such as steel or ammonia production (for fertilizers, not fuel)—is more energy efficient and should thus be given precedence over its use as a fuel.

Hydrogen for heating

Let’s do a simple math. Assuming hydrogen is made using an electrolyser with an energy requirement of 55 kWh/kg H_2 all in, which corresponds to an efficiency of 72%; the hydrogen is then compressed and delivered to the end user with a 90% efficiency, a water-condensing boiler having a 90% efficiency is used for heat application, then the total efficiency is $0.72 * 0.9 * 0.9 = 58\%$. For comparison, we could use a heat pump with an efficiency of 300%, which is 5x higher than H_2 boiler efficiency. Is the choice obvious? This discussion applies to low temperature heating only.

Example of California refueling stations

While our discussion has been mostly theoretical, it is worth looking at real-world data from existing hydrogen refueling stations—such as the 74 public stations in California—to gain valuable insights and to have a more grounded appreciation of the reality. On average, producing 1 kg of hydrogen at these stations uses at least 70–80 kWh, indicating about 40% efficiency only compared to the 68% assumed in our analysis. This energy consumption is much higher than the commonly cited 55 kWh/kg H_2 (excluding compression) due to greater operational losses and unaccounted losses. Although efficiency may improve over time, this large gap highlights the need for realistic expectations about hydrogen production. See: Multi-year energy performance data for an electrolysis-based hydrogen refueling station - ScienceDirect

Conclusion for Part I

This analysis confirms the low conversion efficiency for electricity to hydrogen back to electricity as highlighted by numerous previous studies. However, thermodynamic limitations are not the main limiting factor of hydrogen efficiency in the traditional sense; rather, low efficiency is driven mainly by CAPEX requirements. Additionally, however, the low thermodynamic equilibrium voltage associated with the H_2/H_2O couple puts hydrogen inherently at a 3x disadvantage vs lithium-ion battery competition. This is a physical limitation that further research and development cannot alter. The diffuse state of hydrogen requires additional complexity in handling, causing further losses, which leads us to build a Rude Golberg type system to do the conversion of electricity to electricity; but it is not the main reason making hydrogen conversion inefficient. When all subsystem losses are factored in, the conversion efficiency of electricity to electricity via hydrogen as energy carrier is below 30% for compressed hydrogen and below 25% for liquefied hydrogen. This means that renewable energy infrastructure must be oversized by 4x to support hydrogen applications. This challenge should provide ample motivation for considering an alternative approach.

However, using hydrogen as a chemical feedstock rather than as a fuel can increase efficiency from 30% to 68%, which may support decarbonization in some heavy industries. More detailed application-specific analyses are still needed to fully assess the benefits.

It is important to acknowledge that hydrogen may be suitable for certain specific applications; however, successful deployment depends on targeted system design and integration tailored to the use case. Historically, hydrogen has been treated too much as a generic one-size-fit-all type of solution which has led to a solution that is designed for everything but not optimized for anything. The success of hydrogen deployment may require a holistic system thinking approach right from the conceptual design phase.