International Journal on Advances in Intelligent Systems, vol 16 no 1 & 2, year 2023, http://www.iariajournals.org/intelligent_systems/

Green Hydrogen Production, Transportation, Securitization, and Tokenization using Graph-Neural-Network AI-Automated Control

W. Bernard Lee
HedgeSPA Pte. Ltd.
1 Kay Siang Road #12-02
Singapore, 248922
bernard.lee@hedgespa.com

Anthony G. Constantinides
Financial Signals Processing Lab.
Imperial College London
London, UK SW7 2AZ
a.constantinides@imperial.ac.uk

***Abstract*—HedgeSPA is launching a pilot project for green electricity and hydrogen production in collaboration with a consortium including the Fraunhofer Institute for Reliability and Microintegration (IZM), part of Germany's leading applied R&D institute Fraunhofer Institute. Using IZM's high-temperature and high-pressure sensors designed for the German chemical industry, HedgeSPA will deploy its graph-theoretic AI engine on some of the most advanced hardware supplied (with validation) by Amazon Web Services (AWS) and Nvidia to control the green electricity and hydrogen production processes from the cloud using solar-thermal generation, which can continue generating electricity for up to 12 hours after sunset. While solar-thermal technologies are proven, past deployments required large sites to achieve economies of scale from centralized human support and data centers. Large sites are more difficult to find and as such their regulatory approval processes can be arduous. For the economics of smaller sites to work, they need to minimize or even eliminate human support and data centers from full automation with an AI engine that can adapt to different topographies and respond to weather and environmental challenges using reinforcement learning. In the future, we may further upgrade electricity generation from turbine engines to thermophotovoltaic (TPV) cells, while using the residual heat to superheat water molecules for high-temperature electrolysis, or to Stirling engines (used in conventional-powered submarines built by countries such as Japan and Singapore) with fewer moving parts and no dependency on a water source to produce steam. Today's market has plenty of ready buyers of any green electricity, green hydrogen, or synthetic natural gas produced. The key challenge appears to be financing for the supply side, which can be addressed by securitizing buyers' cash flows to pay for any upfront investments, similar to how homeowners can pay for mortgages instead of rents when an efficient financing mechanism is available to finance the upfront investments for construction companies to build residential homes. We propose to facilitate the financing with a classic waterfall structure under a tokenization mechanism to issue back-to-back tokens against the asset-back securities with high-integrity carbon tokens attached. All of the above fits with a widely-adopted international legal framework for renewable power plants, with only a few new elements that appear manageable based on qualified legal opinions.**

Index Terms*—Graph Neural Network; AI Process Control; Green Energy; Solar-Thermal Generation; Securitization; Tokenization

Note: This paper is the revised and expanded version of [1].

I. Introduction

While commercial photovoltaic (PV) panels have now achieved efficiencies exceeding 20%, typically such efficiency is achieved under optimal radiation conditions [25] [30]. Using PV panels to generate the type of high-voltage, steady-output electricity or hydrogen by electrolysis still faces technical unknowns with low efficiency being the biggest hurdle to wider commercialization [6]. The Ashalim Power Station in Israel [24], the Ivanpah facility in Mojave Desert, California [27], and the Planta Solar 10 in Andalusia, Spain [29], have demonstrated the commercial feasibility of solar-thermal electricity production with efficiencies *starting* from the 20s and reaching the 30s or even more when the plants can operate for up to another 12 hours after sunset [2] [15]. Steam-based turbine generators as well as other heat engines such as Stirling engines are proven, reliable, and cost-effective, and minimize inefficient downtime due to PV panels not being able to operate under optimal solar radiation conditions at all times, especially after sunset.

The project consortium consists of senior researchers in the relevant technical fields, proven professionals in energy project finance, as well as private firms with industrial process control expertise. While typical solar-thermal installations require 150 hectares (or 1,500,000 square meters) to 320 hectares (or 3,200,000 square meters, or 1,789 meters, exceeding 1 mile on each side in a square configuration), we propose to build smaller production sites starting from about 10 hectares (or 100,000 square meters, or 316.2 meters on each side in a square configuration), or at $\frac{1}{32}$ to $\frac{1}{15}$ of typical solar thermal sites today. Such smaller sites are much easier to find and win environmental approvals, and we bring the cost down by controlling a cluster of these production sites on the cloud and servicing them with a technical staff centrally located at one or a couple of locations. More importantly, recent press reports suggest that giant renewable energy plants are at risk from single-point failures that destabilize production and quality [32]. From the buyer's perspective, a single incident at a giant production site may mean writing off millions in anticipated revenues, while the damages from an accident at one plant of a more practical size can be contained. The following is a 8-point summary of this proposed project:

1) *Pilot.* The proposed pilot plant is a ∼5-megawatt facility targeted to meet customary demonstration requirements to secure the form of financing that we intend to pursue. We are aiming for extraction sites with high solar radiation. Green electricity will be either consumed or sent to the power grid where feasible, or it will be used to produce green hydrogen or synthetic natural gas during low-demand hours. Besides traditional turbine engines, we are also exploring other types of heat engines with fewer moving parts or no moving parts to convert heat into electricity.
2) *Production.* The idea is to replicate these extraction sites within reasonable transportation distance from a centralized engineering and transportation hub, with access to fresh, mixed, or seawater, so that, for instance, the produced hydrogen or synthetic natural gas can be shipped to a port for blending with field-produced natural gas for shipment to its final destination. Once proven, this "hub and spoke" production method can be implemented at more than one or two hubs. We will prefer locations with reasonable protection from natural calamities such as hurricanes.
3) *Financing.* The form of financing will be similar to how Tesla uses third-party financing to install photovoltaic panels on roofs across homes in California, where those solar electricity generation units are controlled from the cloud. We are proposing a similar setup because many farms, aquafarms, or ranches can use up to certain portions of unused land for renewable energy without making rezoning applications. Process control will be driven by HedgeSPA's graph-theoretic AI engine working with AWS IOT Core to connect to microcontrollers, as a setup that has been validated by AWS' technical team. This is intended to be a relatively carefree solution for the landowners, who will benefit from access to cheap renewable energy, and potentially steady income streams while complying with government policies of deriving minimum percentages of energy supplies from renewable sources. As standard practice, the annual financing program will cover the cost of construction each year, and the maturity period will be driven by cashflows secured from deliveries.
4) *Buyers.* At least one buyer in Asia has indicated interest to take delivery of up to 2 metric tonnes of hydrogen per day under a highly achievable delivery price target. Other utilities and energy users in Asia have posted tariff schedules on renewables. Today's market challenge is about finding reliable supplies, not about identifying demands for renewable energy. Our assessment of current demand is that we can deliver as much green energy to corporate buyers as we can practically produce, provided that we deliver below a certain reasonable price target. Of course, we must account for any additional processing costs and logistics of getting the end product to buyers. We plan to use high-quality cashflows to support asset-backed securities and attach high-integrity carbon offset tokens to them.
5) *Securitization.* Just like in Tesla's photovoltaic panel use case in California, the smooth financing of the generation units will be the key to the success of the program after our initial pilot project. We plan to finance these projects with asset-backed securities and issue the securities as digital assets to investors, to be custodied using an independently-verifiable ledger. The current legal opinion is that the entire process can be completed within a single jurisdiction if desired. If a third-party jurisdiction is considered desirable, we have discussed a potential issuance with the Hong Kong Securities and Futures Commission, but we are actively exploring other alternatives, such as Dubai, Japan, and Singapore.
6) *Tokenization.* With economies of scale, we are projecting our costs and required investments to be consistent with industry estimates published by Goldman Sachs [26]. To lower the overall costs of financing, we also plan to attach high-integrity carbon offset tokens to these digital assets based on the non-fungible feature of tokens. Given the nature of digital ledgers, with our framework the same one ton of carbon offset credits cannot be sold more than once. Doing so discourages the prevalence of widespread fraud among issuers of carbon offset credits in certain markets today, which often prevents carbon credits in those markets from fetching their fair market values in markets known for high integrity.
7) *Delivery/Shipment.* If the energy is not consumed directly, delivery and shipment can be done using one or more of the following methods: a) Pumping any electricity to the power grid, after using an Automatic Transfer Switch (ATS) to convert to alternating current or by using direct current if the receiving utility can accept direct current, or put the electricity into battery storage first. b) Using green electricity to produce hydrogen for specific applications such as a fuel cell plug-in unit to extend the range for battery-based electrical vehicles. c) Blending either green hydrogen or synthetic natural gas (which can be produced from hydrogen via the proven Sabatier reaction) with field-produced natural gas for power generation. Synthetic natural gas can be injected directly into preexisting liquefied natural gas infrastructure, while shipping hydrogen is a much more complex technical challenge since hydrogen liquefies at the formidable temperature of -253° Celsius and has boil-off rate of at least 1% per day. d) The final option is to deliver hydrogen as ammonia which is a jet fuel/petrol substitute, once a more environmentally friendly process instead of relying exclusively on the current Haber-Bosch process (which is taking more than 1% of the global energy consumption) can be implemented on an industrial scale.

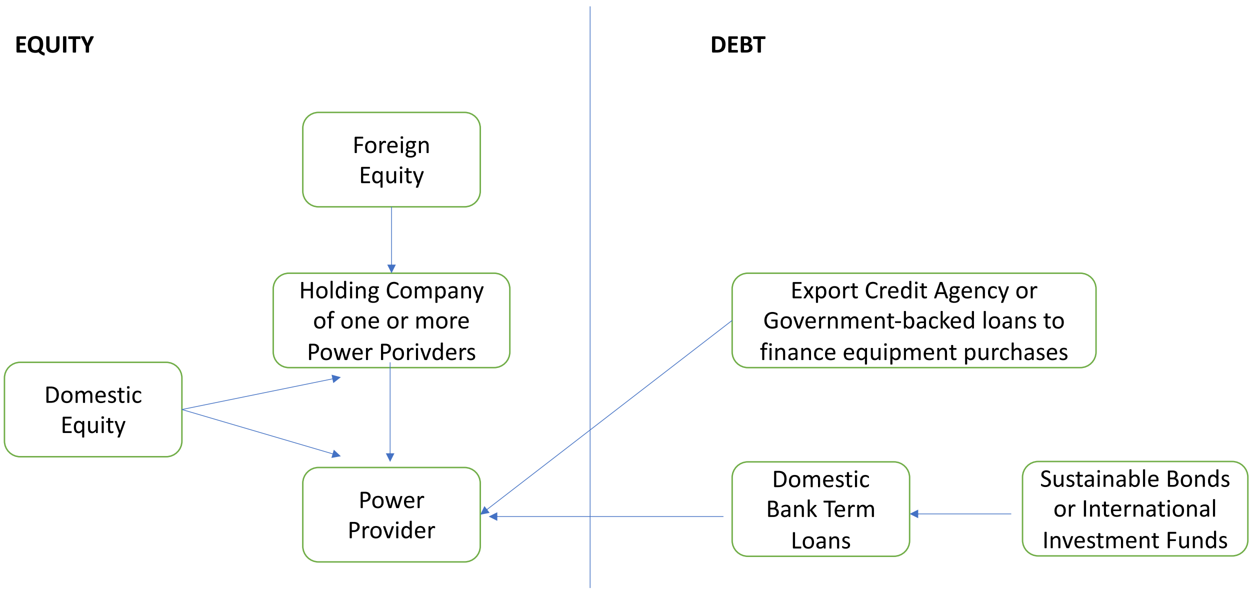

Fig. 1. Typical International Private Power Generation Financing Scheme

8) All of the above fits with a widely-adopted international legal framework for renewable power plants, with only a few new elements that appear manageable based on qualified legal opinions. See Figure 1.

II. Detailed Descriptions

This section provides detailed descriptions of the proposed project:

A. Background Technologies

HedgeSPA has a proven graph-theoretic AI engine with a successful track record in application to investment management and related domains in machine learning [4]. In particular, some of our latest R&D projects have won global recognition by appearing in scientific publications of respected publishers such as Springer-Nature [12]. This proposal aims to deploy this AI engine to control chemical processes for the green energy demonstration plan.

AWS IoT Greengrass device driver can support Web Sockets and MQTT protocols. HedgeSPA's AI engine, already running on AWS, supports Web Sockets, with plans to integrate with MQTT eventually. The planned collaboration model is for IZM to load Greengrass to drive their hardware sensors so that HedgeSPA developers can read the digital outputs from the UNIX ports on sensor microcontrollers. Thereafter, the HedgeSPA technical team can perform the rest of the integration tasks from the microcontrollers to the AI engine.

AWS has also reduced the absolute minimum RAM requirements to 1 megabyte to support light microcontrollers such as the ARM Cortex M-class microcontrollers. Based on the sensor design proposed by IZM, we aim to package the sensors and its microcontroller into a single unit. The microcontrollers can initiate orderly shutdowns in the event that connectivity to the AI engine on AWS is lost due to rare breakdowns, such as peak coronal mass ejections or solar flares which are expected to peak in this decade.

How can this new setup revolutionize the production of green energy? In the past, energy plants had to be large in order to support the cost of a data center to manage the equipment and personnel. This meant that reaching a certain scale was the only financially viable option. However, with the use of an AI engine, we can now create more flexible setups that control the production process from the cloud. For example, commercial farms, aquafarms, and ranches often have unused land that can be leased to build smaller solar-thermal plants, which can then be consolidated at a transportation and maintenance hub for export shipment. By using sensors that can withstand high temperatures and pressure, the AI engine can collect and analyze data more effectively. Additionally, the sensors can be easily relocated within the reaction chambers to optimize AI training without requiring new custom-built chambers, since the sensors are packaged with microprocessors and require only thermalcouple wiring for power supply, which is rated for 1200° Celsius off-the-shelf.

B. Project Objectives

The proposed project's objectives include the following:

1) Reduction in the cost of renewable energy:
   - We expect our hydrogen production cost to eventually get closer to the US Administration's ambitious target of 1 US dollar per kilogram of hydrogen under the economy of scale. The current design is based on the proven technology of Polymer Electrolyte Membrane (PEM) Electrolysis coupled with a Solar-Thermal Steam Turbine and feeding the effluent steam to the electrolysis surface. Both improving the efficiency of the membrane surface using more advanced electrode materials (such as those developed by South Korea at 20 times the efficiency of platinum electrodes) and heating the water (instead of superheating to 2,600 degrees to ionize water molecules into their plasma state under High-Temperature Electrolysis, which may not be energy efficient even when the solar-thermal setup is in a location with very hot weather) will allow a steam turbine or another heat engine to produce the electricity to electrolyze heated water molecules. The pilot plant will allow us to collect experimental data and fine-tune system control parameters to achieve the most efficient temperature/electricity combination given the simulation models and the characteristics of the sites.
   - Active investigation of other safe methods of delivery, such as hydrogenation and dehydrogenation aiming to achieve more practically realizable objectives on such aspects as heat/pressure required in the Methylcyclohexane (MCH) process, including transportation of hydrogen without the significant amount of energy required to liquefy hydrogen at -253° Celsius to transport it in liquid form.
2) Increase in the value delivered by renewable energy:
   - The proposed setup can be augmented by other methods to generate green electricity at non-sunlight hours to maximize the usage of the electrolyzer, which is a fixed-cost investment. For example, there can be a higher attitude wind farm feeding green electricity to the electrolyzers to produce hydrogen at night (slopes cool faster, and cool air rushes down), and all such control mechanisms can be driven by the AI engine on the cloud. This is analogous to how Tesla chooses to sell more expensive electricity to the State of California grid during peak consumption hours during the day but draws cheaper electricity from the public grid to charge vehicles at night.
   - Part of our testing efforts will fine-tune the most efficient temperature-pressure-source combinations based on physical characteristics such as the unique topology of each production site. With such dynamic management, we expect our hydrogen production cost to eventually get closer to the Biden Administration's ambitious target of 1 US dollar per kilogram of hydrogen under the economy of scale.
   - Eventually, our process control engine can consider other factors such as access to consolidated distribution units to optimize production and delivery methods based on the most up-to-date market prices.
3) Improvement in technology readiness and commercial readiness of renewable energy technologies:
   - We start with proven technology and modify the temperate/electricity mix to achieve higher efficiency. This may include using the pilot plant to adapt to lower-temperature techniques tested in laboratories before rolling out such technology in new plants:
     - Sulfur Iodine Cycle
     - Cerium (IV) oxide-cerium (III) oxide cycle
     - Copper-chlorine cycle
     - Hybrid sulfur cycle
     - Iron oxide cycle
     - Zinc-zinc oxide cycle
   - The testing of materials that can stand up to 2,600° Celsius for constructing the reaction chamber in high-temperature electrolysis (likely candidates are ceramics – with some candidate materials already in use in Japan for hybrid-hydrogen power generation), or more efficient materials for PEM electrolysis. Part of our R&D efforts will fine-tune the most efficient temperature-pressure combinations based on the physical characteristics of each production site. There are modeling solutions available for such temperature optimization today [16].
4) Reduction in or removal of barriers to renewable energy uptake:
   - Integration of the AI engine with solar-thermal mirror and control algorithms to focus solar rays on the reaction chamber and with actual testing of the thermal efficiency on the demonstration plant means that any control software updates can be deployed directly from the cloud. While we expect to staff support professionals servicing multiple sites around each major delivery center to meet health and safety requirements, these plants should no longer require large data centers and large crews to run. Just like Tesla in California, we are proposing a smart solution that can be deployed flexibly by any farm or aquafarm owner with a low barrier to entry.
   - The proposed computational control architecture has already been validated by AWS to work with AWS IOT Core, while the AI engine is already working with WebSockets as required by AWS. Deploying the control mechanism on the cloud will mean that many farms or aquafarms with at least 5$\sim$10% of unused land will have access to cheap renewable

energy as well as income streams while contributing to the global fight against climate change by delivering green energy to a transportation hub for domestic consumption or exporting.
- Finally, one key to success in the world's sustainable energy transformation is the availability of financing. We are proposing an innovative way to finance these projects using digital tokens based on the fact that ample demands and supplies exist, but the upfront investment simply needs to be made.

5) Increased skills, capacity, and knowledge relevant to renewable energy technologies:
   - Physical testing of all chemical steps to fine-tune the chemical simulations already completed on the Chem-CAD software ASPEN – In particular, future upgrades will follow the path of simulation to lab-scale testing to demonstration-scale testing. Where feasible, the pilot plant can be used to test some of the lower temperature cycles in electrolysis. This will also involve the future construction of lab-scale demonstration units that will allow a continuous flow of chemical processes rather than the hitherto batch mode operations. These activities will include the practical training of scientists, engineers, and technicians.
   - The world cannot migrate away from fossil fuels based on batteries alone because the use of batteries in commercial transportation such as buses and trucks has the known problem of low energy-to-weight ratio with the need to transport the heavy weight of batteries. In fact, batteries finish dead last in an energy density chart published by the top American manufacturer of locomotives. In the same chart, the manufacturer suggests that hydrogen can be one solution to this problem. Hydrogen fuel cell supplement is one attractive solution for mid-distance commercial vehicles with its far more attractive energy-to-weight ratio, but voltage stabilization and heat mitigation are necessary to avoid the overheating of batteries and therefore potential explosions. Moreover, HICE (hydrogen internal combustion engine) or the practice alternative of mixing fossil fuels with ammonia requires no to extremely simple modifications to the existing gasoline-powered vehicles, e.g. HICE has been deployed on buses in California for over a decade. As we need to target our capacities based on end customer demand, we will also devote resources either to test the downstream technology directly or work collaboratively with relevant R&D institutes.

III. Proposal Implementation

The generation of hydrogen via electrolysis is well understood by scientists. We will start with a proven setup and explore other alternatives to improve efficiency and lower reactor temperature:

A. Proposed Hydrogen Production Technology

The typical consumption rate of hydrogen by a passenger vehicle has been estimated to be slightly above 10 grams/kilometer. Assuming that commercial vehicles travel 200 kilometers/day, the average daily consumption of each vehicle will be $\sim$2 kilograms of hydrogen. One source of hydrogen demands that we have identified is from a fleet of 1,000 vehicles, which will consume 2 metric tons per day or about 700 metric tons/year (assuming 350 operating days/year, allowing for downtime due to public holidays and maintenance). The demand volume is expected to increase dramatically over the next few years. We expect the aggregate demand that we can serve can eventually exceed the yield of 2,000 metric tons/day.

We propose to produce hydrogen in scale using solar-thermal electrolysis, which can be either based on a) High-Temperature Electrolysis or one of its lower-temperature variants based on a different chemistry cycle, b) Polymer Electrolyte Membrane (PEM) Electrolysis coupled with a Steam Turbine or a Stirling Engine (a heat engine design that does not require steam, see Figure 3) and feeding the effluent steam to the electrolysis surface, as illustrated in Figure 4. Both modifying the electrode membrane and heating the water will allow a steam turbine to produce electricity to electrolyze vaporized water molecules. The high-level design is to produce steam in a solar-thermal plant, convert the thermal energy into electricity, and then use the effluent to feed steam to the electrolyzer surface to produce hydrogen. Part of our R&D efforts will fine-tune the most efficient temperature-pressure combinations based on the physical characteristics of each production site. There are many simulation models available for temperature optimization [16].

Nature reported that MIT achieved 40% efficiency (comparable to turbines) with a heat engine with no moving parts by reflecting any unconverted photons back to the emitter. This technology is known as Thermal Photovoltaic (TPV) cell [11]. While the high efficiency was gained by preserving heat, the process of conversion may still be slow especially if it must be flow-based or the heat circuit may overheat. Practical application is still a challenge because ceramic materials such as zirconium must be used to construct a heating/reaction chamber without the risk of its melting and then connected to a "cooling" coil with Gallium Arsenide (GaAs) multi-junction devices and back reflectors to convert infrared radiation into electricity. The coil can be filled with molten salt or a heat transfer fluid, and then reheated. Any residual heat can be used to plasmarize water molecules for high-temp electrolysis. In summary, the building up and maintaining of such high temperatures for efficient conversion do present practical challenges to developing a practical heat engine with no moving

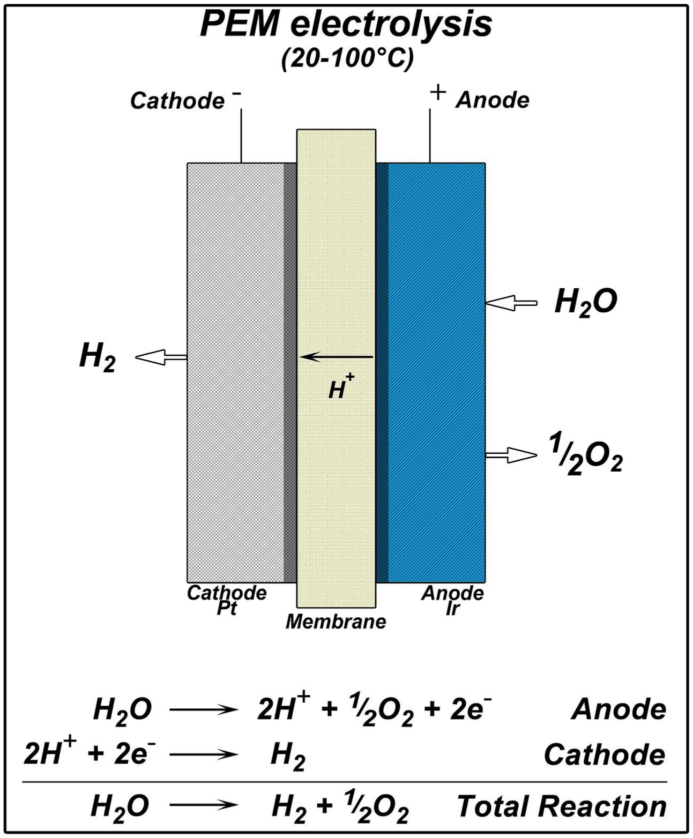

Fig. 2. PEM Electrolysis Process (Davidlfritz, CC BY-SA 3.0 <https://creativecommons.org/licenses/by-sa/3.0>, via Wikimedia Commons.)

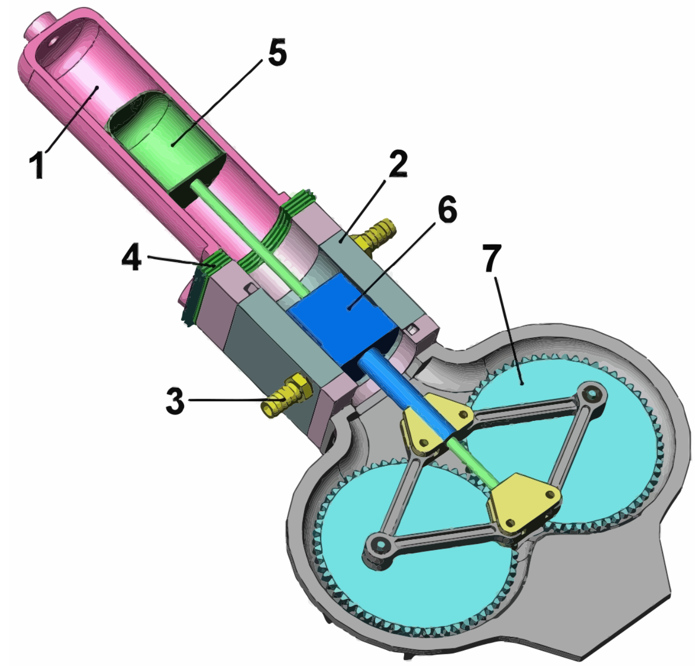

Fig. 3. Cut-away diagram of a rhombic drive beta configuration Stirling engine design: 1 - Hot cylinder wall; 2 - Cold cylinder wall; 3 - Coolant inlet and outlet pipes; 4 - Thermal insulation separating the two cylinder ends; 5 - Displacer piston; 6 - Power piston; 7 - Linkage crank and flywheels Not shown: Heat source and heat sinks. In this design, the displacer piston is constructed without a purpose-built regenerator. (BetaStirlingTG4web.jpg: Togoderivative work: Ionutzmovie, CC BY-SA 3.0 <http://creativecommons.org/licenses/by-sa/3.0/>, via Wikimedia Commons).

parts or deployable thermal energy grid storage (TEGS). The high cost of GaAs fabrication is also a potential roadblock but it has been reduced thanks to new fabrication techniques. This can be an interesting supplement to the next phase of our R&D program after starting with steam turbine or Stirling generators.[1]

For solar-thermal generation, the most desirable region to run such productions will be hot and dry desert areas with over 3,000 sun-hours per year and access to water. Any location with over 2,000 sun-hours per year can still be a reasonable location for demonstration purposes. The steam will be extracted by solar-thermal distillation and the high temperature will be generated by AI-controlled solar mirrors.

The currently proposed design will be based on a triple-heat-circuit solar-thermal design, as shown in Figure 4. TPV cells can supplement and eventually replace the turbine engine in the heat circuit.

In the future, besides TPV cells, such a setup can be augmented by wind, photovoltaic or tidal methods to generate green electricity, and with the following lower temperature chemical cycles than high-temperature electrolysis to produce hydrogen:

- Sulfur Iodine Cycle
- Cerium (IV) oxide-cerium (III) oxide cycle
- Copper-chlorine cycle
- Hybrid sulfur cycle
- Iron oxide cycle
- Zinc-zinc oxide cycle

B. Delivery

Most sites will be of dual use. Feeding electricity to a power grid involves the proven technology of using an Automatic Transfer Switch. When hydrogen is produced, the harder technical challenge is the shipping of hydrogen to its final delivery sites, which can be complex and dangerous. When the cost of and access to land is not an issue (e.g. Australia), then we should produce the renewable energy as inexpensively as possible and then create the hydrogen. Then, one can afford a more expensive method to ship the green energy over a longer distance. On the other hand, when the cost of and access to land is an issue (e.g. an industrial park in Asia to supply energy next door, where land is allocated), then we can focus on creating maximum yield from whatever land available and using the electrolyzer only when there is no immediate demand for the green electricity. The possible methods of transporting hydrogen are described as follows:

[1]The list of GaAs foundries with the capacities to launch any volume production of TPV cells is available from https://en.wikipedia.org/wiki/List_of_semiconductor_fabrication_plants.

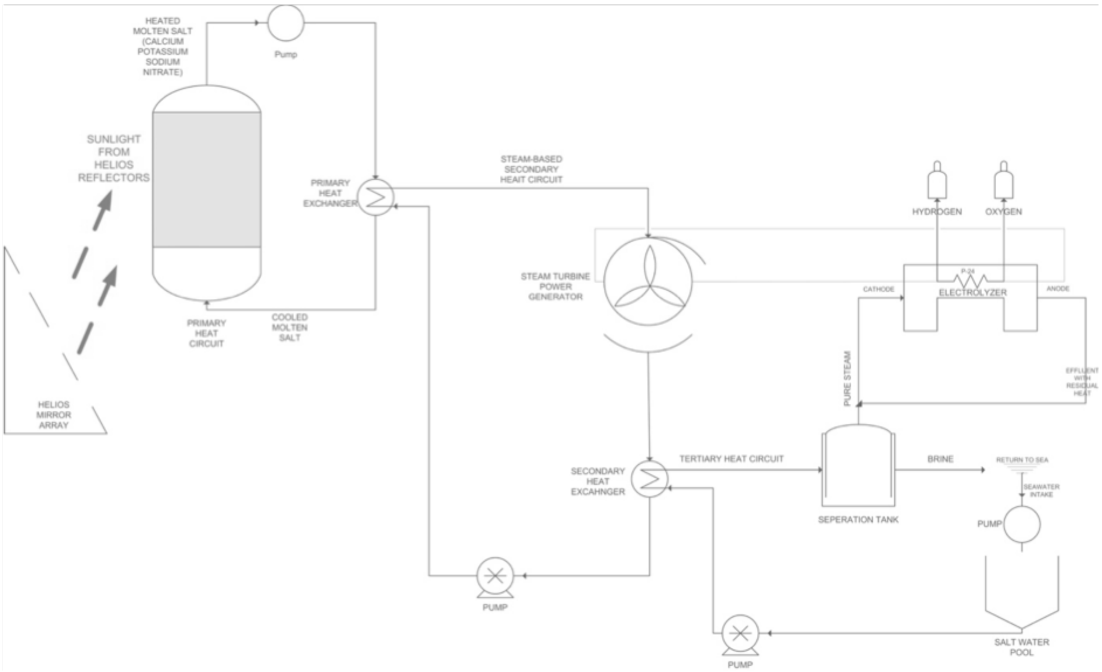

Fig. 4. Solar-Therma Hydrogen Production Plant Schematic

1) For shorter distances (within 3,000 kilometers, we propose that we ship with Type 3 or Type 4 compressed hydrogen tanks. This is most likely the solution when the primary goal of a renewable energy plant is to feed green electricity to an electric grid. Hydrogen is produced to capture the surplus green energy, and shipped to other flexible demand buyers such as hydrogen fuel cell distributors. As an example, this is a good solution when limited land is allocated adjacent to an industrial site, with the goal of delivering green energy for direct and immediate consumption.
2) For mid-distance deliveries of about 3,000∼7,000 kilometers, one solution is to use bulk carriers to ship methylcyclohexane (MCH). This method involves the hydrogenation of toluene, a colorless and odorless liquid, into methylcyclohexane. After arriving at a receiver port, the MCH can be dehydrogenated (under heat and pressure) to release the hydrogen for local consumption. Currently, hydrogen delivered with this method can be delivered to locations such as the Port of Kawasaki in Japan. This is seen as a temporary workaround because unused refineries can be converted into MCH hydrogenation and dehydrogenation plants, thus lowering the cost of the remewable energy supply chain using this solution.
3) For delivery over longer distances, we propose to convert the hydrogen into synthetic natural gas using the Sabatier reaction first and then ship the liquefied natural gas (LNG). Some utilities in Asia have been given explicit targets to blend synthetic natural gas with LNG by 2030. Hydrogen has a liquefication point of -253° Celsius, while LNG liquefies at -161.5° Celsius at atmospheric pressure. Moreover, liquid hydrogen has a boil-off rate as high as 1% per day, meaning that 20 $\sim$ 30% of a shipment may be dissipated before reaching a long-distant destination. Synthetic natural gas is a far more practical method to ship to long-distant locations of over 10,000 kilometers. This is the likely solution when there is an existing LNG infrastructure and plenty of land for generating renewable energy, such as the Middle East or parts of Malaysia, with plenty of carbon feedstock from their natural gas fields.
4) The problem with synthetic natural gas is that the carbon feedstock will still be released into the atmosphere. To address such a problem, we propose to use ammonia as a subsequent-stage solution for transportation fuels, since it liquefies at the easily achievable temperature of -33° Celsius. Ammonia can be deployed as a direct substitute for gasoline and jet fuel in certain applications, often requiring no to relatively simple modifications to the power plant. Upon combustion, ammonia emits nitrogen and water. Neither of them causes any harm to the environment. Unlike MCH which requires the shipment of toluene back to the production site, ammonia does not require any dehydrogenation at the delivery point and the "return shipment" of toluene, resulting in more cost savings.

   Today, the industrial conversion of hydrogen to ammonia is based on the extensively-used Haber-Bosch process. The mixture of nitrogen and hydrogen will be compressed to 100 ATM but cooled to 450° Celsius using a cooling coil. The resulting gas will be fed into an absorber column filled with water. The ammonia is extracted from the bottom of the absorber column. Any unreacted gas will be recycled from the top of the absorber column. This needs to be done as a batch-based process because of the extreme pressure required. A more pragmatic strategy in the future is to use a bio-catalyst, such as using nitrogenase to combine hydrogen with nitrogen to make ammonia, which is significantly ”greener” than the Haber-Bosch process with a much lower excitation energy threshold. In other words, hydrogen is pumped into holding tanks containing organisms with the enzymes to convert hydrogen into ammonia. The issue is how to do so beyond lab-scale experiments given that ammonia is highly solvable in water.

Common methods of removing ammonia from the output solution either oxidizes the ammonia (thus defeating the purpose) or uses heat and pressure, leading some to question how green the produced ammonia is. Nonetheless, ammonia is a possible solution when there is an abundance of land for generating cheap renewable energy, and therefore the commercial objective is to make the renewable energy as cheaply as possible from locations such as North Africa, the Middle East or Australia. Then, it is a matter of turning green hydrogen into green ammonia, even by using a relatively slow ammonia removal pathway that requires large holding tanks, to be delivered for use in say commercial shipping, where the ammonia is burnt so a tiny percentage of impurities is considered acceptable. When there is sufficient volume, the high minimum cost of building a Haber-Bosch plant, to be eventually replaced by biochemical processes with lower excitation energy thresholds, can become commercially justifiable.

C. Securitization and Tokenization

This subsection contains the implementation details of the securitization and tokenization piece of the project, to create the so-called “high integrity” carbon offset tokens.

1) Token-based Financing Use Case:

1) Investors can buy digital assets on renewal energy using an authorized central bank digital currency (CBDC) or a stablecoin backed by a fiat currency.
2) Their local regulators can confirm the eligibility of the investors to invest (instead of relying on payment apps or commercial banks to do KYC/AML) according to local regulations.
3) The identity of the qualified buyer can be lodged in a “custody” ledger at a brand-name custodian that will be kept under the watchful eyes of the issuance jurisdiction’s central bank or securities regulator.
4) The issuance does not need to be for any specific jurisdiction such as Hong Kong. We are using Hong Kong as one example because of a Green Bond program already in existence, but we may choose to do so under the CBI program in Australia or any other relevant jurisdiction such as Singapore, in case Hong Kong may require complicated licensing for issuing digital tokens.
5) If the security is redeemed (either at maturity or by the market maker), the proceeds can be converted back to the CBDC based on the prevailing exchange rate.
6) If working with one specific regulator may take too long, then a similar idea can be applied to another jurisdiction, e.g. Japanese investors provided that the tokens comply with ERC-721 requirements. We have been working with one of Japan’s top two securities law firms to determine how we can meet the exempt status for selling to Japanese investors. For a simpler proof of concept without excessive regulatory concerns, another possibility is to issue the paper securities while running the digital token ledger in the same jurisdiction.
7) The HedgeSPA core investment platform will integrate with Hyperledger Fabric and offer the users of our solution a secure way to manage both the tokens and the underlying carbon assets. Doing this should create no unwarranted security concern from investors in that the actual ledger sits with a brand-name custodian.
8) We also intend to move some of our non-asset-backed securities investment products to such a ledger to create digital ETFs. This way, we attract more traditional investors to use digital tokens as an asset class for financing different sustainable energy projects. The more precise technical details are given below and further illustrated in Figure 8:
   - Construct an investor portal[2] for an investor to transact on one unit of security.
   - API calls working with CBDC portal to buy/subscribe to the asset-backed security:
     - investor transfers the quoted CBDC amount to the CBDC account of the custody bank
     - local regulator allows the custody bank to confirm investor eligibility (preferably electronically);
     - log one unit of the security under the investor’s name in the bank ledger;
     - custody bank to issue one token (on the Hyperledger Fabric blockchain) to the investor.
   - API calls working with investor portal to sell/redeem asset-backed security:
     - investor to give instructions to the custody bank to sell one unit of the token on the portal;
     - confirm investor instruction to sell based on the latest bid price quoted by the market maker;
     - confirm account details and double-check that the account has no KYC issues (based on local regulatory requirements);
     - sell or redeem one unit of security; reverse one unit of the security under the investor’s name in

[2]The initial version can be based on a subset of our fintech terminal interfacing with the API of a digital market such as Coinbase, with a subsequent version upgraded to API calls to an in-house marketplace running the Hyperledger Fabric chaincode.

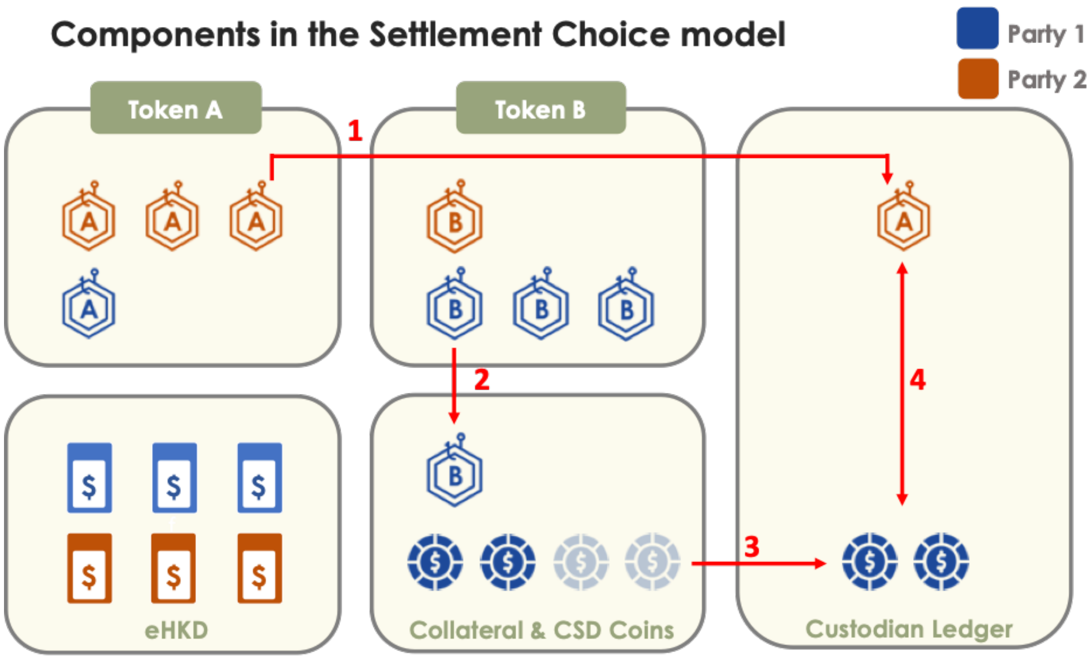

Fig. 8. Asset Tokens in Custodian Ledger

the bank ledger;
  - custody bank to deposit proceeds to the investor's CBDC account.
- API calls to detach carbon offset token from combined security:
  - investor to surrender token on the portal and give instructions to detach carbon token from security[3];
  - reverse one unit of the combined security under the investor's name in the bank ledger;
  - add one unit of the standalone security and one unit of carbon credit token under the investor's name in the bank ledger;
  - sends one unit of the standalone security and one unit of carbon credit back to the investor's account.

2) Benefits from Token Financing Proposal:

1) Financial inclusion – CBDC does not receive any interest payment unless it is deposited. Today, inflation is driven heavily by surging energy prices, and the world is getting ready for its expected transformation to adopt renewable energy. The creation of straightforward ways to convert CBDCs into renewable energy-linked securities allows the financially "less involved" population to participate in the upside of adopting renewable energy.
2) Minimize fraud and develop high-integrity carbon market in Asia – The regional carbon credit market in Asia has not been able to fetch international carbon prices to date due to the (mis)perception of widespread fraud. Attaching non-fungible serial numbers to digitalized carbon credits is one solution toward developing a high-integrity carbon market since the same one ton of verified carbon credits on the blockchain can no longer be sold more than once. Selling the same one ton of carbon offset credits more than once is how certain Asian carbon markets has lost the trust among international investors.
3) A mechanism to settle cross-border transactions in securities issued in one jurisdiction for other CBDC and stablecoin investors – This is expected to become a highly credible test case for interoperability with large-value transactions under delivery-versus-payment (DvP) in a shorter settlement timeframe than t+2 via a brand-named custodian, in the following two ways:
   - Settlement on a programmable CBDC can be made as soon as a digital asset is received and the ledger entries are updated. There is no other reason for this transaction to require T+2 settlement. As and when both the custodian and the investor agree, this can be settled based on even end-of-day or T+0 settlement. Please note that in a typical setup like this, the paper security will be held by the custodian in "street name". This mechanism should be done inside the custodian and does not require any designated "oracle" unless required by regulators.
   - Similarly, a digital bond can release its coupons as programmable CBDC automatically when certain pre-defined conditions (identified by a designated oracle) are met. Without the cost of servicing agents to send out coupon checks, digital bonds can be easy

[3]Once detached investors cannot reattach based on industry practice, similar to detaching a warrant.

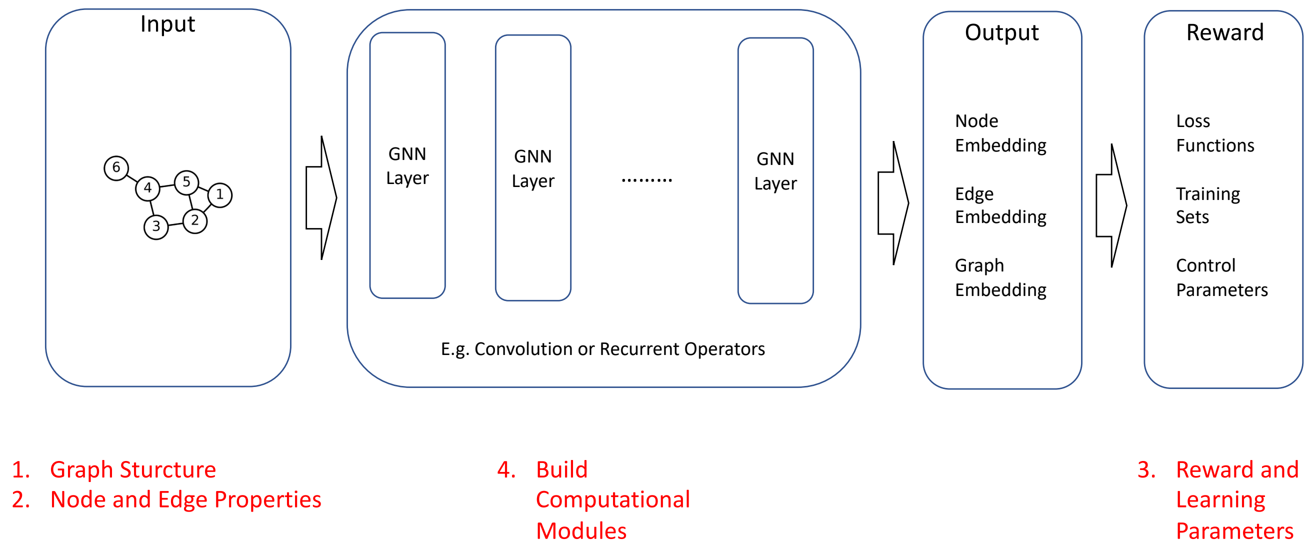

Fig. 9. Typical Steps in Designing a GNN

to manage to accommodate more retail investors, especially retirees. Today, one reason for the high minimum on most bond investments is that issuers are unwilling to pay for the high costs of servicing retail investors (e.g. sending coupons, and manual corrections of common book entry mistakes).

3) Initial Use Case: For our initial test case, we have attracted the interest from one financial institution in Asia which has secured certain ability to issue digital wallets for their local investors. We are in discussion with that institution because our tokens may fit their investors' interest in more diversified investments and are safer than equities. This may be an effective test case by offering digital securities directly to local investors without the explicit requirement to involve a third-party jurisdiction. Simultaneously, we are also in discussion with other potential issuance parties in that we are interested in identifying a general solution that is not limited to any specific country, e.g. one country offers potential test sites, on-site green energy demands as well as the off-shore exchange to issue the asset-backed tokens. We are also aware of that one major technology giant is actively involved in the development of Non-Fungible Tokens. We are agnostic about which base technology to use to develop the necessary blockchain mechanism. It is primarily about doing so in the most cost-effective way without creating strong dependency on a single technology/provider.

IV. Control Mechanism by AI Engine

In this section, we present our proposed method for production control using an AI engine based on graph theory, which has been developed and proven by HedgeSPA. The approach involves creating a graph model with nodes and edges that represent M stages, N processes, and $M-1$ holding tanks involved in the production process. This type of model has been previously used by organizations like NASA for automated control applications, and has been proposed by various sources, including [9], [7], and [21]. To explain our approach clearly, we break it down into four steps, which are depicted in Figure 9.

The first step involves modifying the framework presented in [9] for manufacturing automation to create the graph structure that accommodates the complexity of chemical processes. The second step involves specifying the properties of nodes and edges in the graph structure using Graph Neural Network, with nodes representing reaction chambers and links representing fluid flows. In the third step, we design reward and learning parameters, including real-time sensor states and system-wide loss functions, and apply smoothing functions to increase the robustness of estimates. Finally, in the fourth step, we build computational modules using the graph abstraction of the green energy production process, integrating the HedgeSPA graph-theoretic AI engine with IZM sensors for optimal control. We apply multiple computational methods, such as Multi-Agent Reinforcement Learning (MARL) and spectral techniques based on graph Fourier transforms, to solve for the optimal solution.

A. Graph Neural Network (GNN)

The proposed approach utilizes a standard graph notation, where $G = (V, E)$, with V representing the set of nodes $x_i \in X$ and E representing the set of links between nodes $e_{ij} = (x_i, x_j) \in E$. The adjacency matrix A is defined as an $n \times n$ matrix where $A_{ij} = 1$ if nodes v_i and v_j are connected, represented by $e_{ij} \in E$, and $A_{ij} = 0$ if they are unconnected, represented by $e_{ij} \notin E$. The hidden states and output vector of a node v are denoted by $\mathbf{h}_v$ and $\mathbf{o}_v$ respectively, following standard graph-based learning notation.

To represent a node, the proposed approach uses a two-step process under the spatial method, as suggested by [9]. In the *Aggregate* step, the hidden states of immediate neighbors of node i are aggregated, i.e. $a_i^l = Aggregate\left(h_j^{l-1} : j \in \mathscr{N}(i)\right)$. In the *Combine* step, the node embedding for the l-th GNN layer is generated, i.e. $h_i^l = Combine\left(h_i^{l-1}, a_i^l\right)$. The node embedding is represented by h_i^l, where l denotes the number of layers of aggregation, and $\mathscr{N}(i)$ denotes the set of immediate neighbors of node x_i. The proposed approach allows for l layers of aggregating the node embedding, while weight matrices are not explicitly used in the Combine step, according to [9].

Different standardized notations have been suggested by other researchers across multiple training methods, as discussed in survey and review publications on GNN. However, the notation used by [9] is primarily followed in the approach proposed here to allow for easy comparison and contrast between different types of applications, and to ensure consistency in variable indexing among developers. It's important to note that the definitions of the graph, nodes, and edges described in the subsequent subsections are specific to the chemical engineering problem being addressed, rather than the manufacturing problem presented by [9].

B. Production System-Level Modeling

This subsection explains how the production system is modeled to account for the interactions between the different stages and processes involved. The system has multiple stages, each with one or more reaction chambers, as well as intermediate holding tanks between each stage. The level of each holding tank is tracked, and the capacity of each tank is limited.

First, the system is assumed to have M stages, N processes, and $M-1$ holding tanks. Each reaction chamber in the system processes one step of the chemical process. However, one stage may have one or more reaction chambers before feeding into a holding tank. A quality inspection is carried out in the final stage of the full chain of chemical reactions.

Intermediate holding tanks are present between every two consecutive stages because each stage may have varying throughputs during different stages of its chemical reactions. The level of the m-th holding tank is denoted by b_m, where $n \in [2, ..., M]$. The holding tanks have finite capacities, with B_m representing the capacity of the m-th holding tank. Thus, $b_m \leq B_m$.

If a reaction chamber is operational but its upstream holding tank is empty, it is considered "starved", which can cause downstream delays and disruptions. If a reaction chamber is operational but its downstream holding tank is full, it is considered "blocked", which can also cause delays and disruptions upstream. These conditions are undesirable and can significantly reduce system yield if not managed properly.

The decisions made regarding process control and the conditions of each reaction chamber directly affect the system-level dynamics of the production system. The resulting output can be highly non-linear and characterized by step-function-like behavior, highlighting the importance of effective management and control of the system.

C. Reaction Chamber-Level Modeling

This subsection explains how the sensor states and observation uncertainties are modeled at the reaction chamber level, which is essential for making decisions at both the system-level and process-level. To capture the sensor state, discrete values are used as sensors send digital outputs. A scalar representation is used to simplify notations, although multiple sensors attached to a single reaction chamber can be treated as a vector.

The sensor state of process n after processing k chemical steps is denoted as $x_{n,k}^t$, where $\mathscr{T}$ is the maximum number of possible sensor states, and the superscript t is used to identify a specific sensor. The observation of the sensor state at the k-th pass is denoted as $\tilde{z}_{n,k}^t \in \mathscr{Z} \equiv [1, 2, ..., \mathscr{T}]$. The superscript t denotes a specific sensor. The transition model of the sensor state is the probability of the sensor state transitioning from state $x_{n,k-1}^t$ at the $(k-1)$-th pass to state $x_{n,k}^t$ at the k-th pass, and can be written as the conditional probability $p(x_{n,k}^t | x_{n,k-1}^t)$, and the observation model of the sensor state is the probability of observing $z_{n,k}^t$ given the underlying sensor state is $x_{n,k}^t$, or the conditional probability $p(z_{n,k}^t | x_{n,k}^t)$.

The transition model can be derived through chemical engineering modeling or experimentation to account for practical losses and sensor delays. The observation model, which is independent of specific sensor placements or machine learning models, can be obtained from sensor data calibration and machine learning error analysis.

D. Process-Level Modeling

This subsection discusses the process-level modeling which involves creating desired compounds from raw materials using chemical engineering processes. As each chemical process can be unique, it may not be practical to include all processes and constraints in a single graph. To address this, a general mathematical model is presented that focuses on multi-level integrated control.

First, a chemical substance is characterized by its key properties, which are denoted as q_m^i, where m represents the stage of the process ranging from 1 to M. These properties can be represented as a vector.

The process model is represented by the equation $q_m^i = p(q_{m-1}^i, x_{n,k}^t, u_n^i, \epsilon)$. Here, q_{m-1}^i is the chemical or physical feature before the current stage, u_n^i is the process control parameters applied to the chemical, $x_{n,k}^t$ is the sensor state of the reaction chamber at the current time, and ϵ represents random factors within the process.

The final feature of the chemical after the last stage, M, is q_M^i. To ensure quality, the final feature q_M^i is compared with a quality standard Q^*, such as purity. If q_M^i is within Q^*, the chemical is compliant with the qualify standard; otherwise, it is non-compliant.

In summary, the process model describes how a chemical's properties change from one stage to another, depending on the process control parameters u_n^i and sensor state $x_{n,k}^t$. The final feature of the chemical is compared with a quality standard to ensure it meets the required quality.

E. Objective Function of Overall System

This subsection explains the objective function of the overall system, which aims to find the optimal sequence of control actions for each process and reaction chamber to maximize the expected production yield of the system. The objective is to solve for the maximum value of the expected production yield, denoted by π^*, by adjusting the process parameters while considering the system constraints, denoted by C. The total production yield during a specific time horizon, $y(T;\pi)$, is affected by the control actions, u_n^i, which follow the control action sequence, π. In mathematical form, the objective function can be expressed $\pi^* = \underset{\pi}{argmax} \mathbb{E}_\pi[y(T;\pi)]$, subject to C. The system constraints include the chemical and physical properties and limitations, as well as the system, reaction chamber, and process-level models described in the previous subsections.

F. Graph Modeling of the Production System

In order to model the production system, a graph is constructed where each node represents a reaction chamber, and the edges represent the fluid flows between them. The integrated control problem requires that each node is characterized by a set of features denoted as $x_n = [b_m, B_{m+1} - b_{m+1}, x_{n,k}^t, w_n, \alpha_n, q_m^i, u_n, n, m]$. These features include various parameters, such as the immediate upstream holding tank level (b_m) and the downstream holding tank vacancy ($B_{m+1} - b_{m+1}$), which is used instead of the holding tank level to account for any blockages caused by the downstream holding tank.

Other features include the sensor state of the reaction chamber ($x_{n,k}^t$), the operating status of the reaction chamber (w_n and α_n), the chemical or physical characteristics of the chemical being processed by the reaction chamber (q_m^i), and the current process control parameters applied to the reaction chamber (u_n^i), among others.

In addition, each reaction chamber is uniquely identified by the parameters n and m, which facilitate the sharing of parameters in the neural networks. It's worth noting that the proposed computational framework is based on a spatial technique, but its extension to spectral techniques will be discussed further in subsection H.

G. Smoothing Functions for Sensors

There are various established techniques for improving and smoothing the signals obtained from electronic sensors, such as exponential smoothing and Recursive Bayesian Estimation (RBE), mentioned in [9]. Exponential smoothing eliminates noise and local shocks from the raw signals, while RBE is a more advanced method for recursively estimating the sensor state and obtaining our belief about the true sensor state.

In the RBE process, the sensor state belief at the previous $(k-1)$-th pass is used to predict the current sensor state at the k-th pass using the transition model defined in subsection C: $p(x_{n,k}^t | \tilde{z}_{n,1:k-1}^t) = \sum_{x_{n,k-1}^t=1}^{\mathscr{T}} p(x_{n,k}^t | x_{n,k-1}) p(x_{n,k-1}^t | \tilde{z}_{n,1:k-1}^t)$ We can update this sensor state belief using the observed sensor state and our knowledge of observation uncertainty, such as cutting off negative temperature readings at 0° Celsius or ambient temperature.

To obtain the embedding of node n, we set $h_n^0 = x_n$ and use the updated Aggregate and Combine rule $h_n^{l+1} = \sigma\left(\sum_{j \in (n' \cup n)} \alpha_{nj}^l h_j^l W_V^l\right)$. Here, $\sigma(*)$ is an activation function, n' denotes all immediate neighboring nodes to node n, W_V^l is a learnable weight matrix applied to the embedding h_j^l from the previous layer, and α_{nj}^l is the weight assigned to node j. This process is similar to exponential smoothing and other signal processing filters but can incorporate domain knowledge of observation uncertainty based on the chemical processes.

H. Computatonal Model for Learning

To efficiently train any GNN to achieve the theoretical yield of a chemical process, an effective computational model must be created. This can be achieved through supervised or semi-supervised learning, with labeled data mixed with unknown or unlabeled components representing deviations from theoretical yields. The MARL algorithm proposed in [9] can be used to optimize the learnable parameters of the GNN and complete the training process.

The MARL approach is advantageous over a single-agent scheme for two main reasons. Firstly, the combined actionable space for the system grows exponentially with the number of reaction chambers, making the problem computationally intractable if formulated as single-agent reinforcement learning. Secondly, non-linearity often arises from blockages and starvations, which can be managed at the individual agent level using penalty functions triggered by blockages and starvations.

The control actions in MARL are responsible for setting the appropriate process control parameters for each reaction chamber represented by an agent. The control action sequence is denoted as $\pi(u^n | h_n^L)$, where L is the number of layers in the GNN and h_n^L is the final node embedding output from the GNN. The reward function evaluates the effectiveness of the joint action given the system state and drives the agents to improve their control policy actions, aiming to improve system-level yield. The reward function is generally defined as $r_t = y_t - d_t$, where y_t is the discretized production yield, and d_t is the discretized penalty function used to measure non-compliance of the final output with the required quality standard. The reward function aligns with the problem objective of maximizing production yield while minimizing non-compliance with required quality standards.

In summary, implementing the MARL algorithm is a reliable computational model for learning the theoretical yield of a chemical process. This involves defining control actions, discretizing the continuous timeline and chemical volume, defining the policy and reward functions, and optimizing the learnable parameters of the GNN to improve the system-wide expected production yield.

I. Computational Framework for Training

To train the spatial neural network described earlier, the control policy action gradient for each agent can be optimized. This gradient is defined by $\nabla J = \mathbb{E}_\pi \left[\sum_n \nabla \log \pi(u^n|h_n^L)(Q(s,\mathbf{u}) - V_{tot}(s))\right]$, where $Q(s,\mathbf{u}) - V_{tot}(s)$ is the advantage function associated with the neural network's learnable parameters.

To differentiate the learnable parameters for each agent during the training process in a multi-agent spatial framework, the conjugate gradient method can be used. However, if J is highly nonlinear, a gradient-based solution can be inefficient and inaccurate due to the tendency of the algorithm to circle around local optima.

J. Spectral Techniques and Other Training Methods

The previous subsection explained a computational framework for training to achieve the theoretical maximum yield of the control process, using a process similar to that shown in Figure 9. This framework can be retrained as needed to get closer to the theoretical maximum yield. The multi-agent reinforcement learning (MARL) algorithm is advantageous for achieving an optimal solution to this problem because its structure is generally consistent with the problem's structure. Other possible computational modules for training include sampling and pooling operators, and several open-source implementations of GNN solvers are available, such as PyTorch, MPNN, and GraphRNN.[4]

Although spatial techniques have their advantages, they are limited because they are variants of the conjugate gradient method and are often restricted to working with the first and at most second-order derivatives. Therefore, other techniques, such as spectral techniques, may be useful for stabilizing or speeding up specific steps in the learning process. Spectral techniques are based on a spectral representation of the graph Laplacian $\mathbf{L} = \mathbf{D} - \mathbf{A}$, computed from the degree matrix $\mathbf{D}$ and adjacency matrix $\mathbf{A}$ of the graph. They are interesting from a computational efficiency point of view, especially when the system graph gets bigger and more complicated from the use of multiple sensors for each reaction chamber. In essence, a graph signal $\mathbf{x}$ is firstly transformed to the spectral domain by the graph Fourier transform $\mathscr{F}(\mathbf{x}) = \mathbf{U}^T\mathbf{x}$, then the convolution operation is conducted. After the convolution, the resulting signal is transformed back using the inverse graph Fourier transform $\mathscr{F}^{-1}(\mathbf{x}) = \mathbf{U}\mathbf{x}$. $\mathbf{U}$ is the matrix of normalized eigenvectors on the graph Laplacian $\mathbf{L} = \mathbf{U}\mathbf{\Lambda}\mathbf{U}^T$ of the system graph. The graph Fourier transform is based on the eigenfunctions of the one-dimensional Laplace operator and was proposed by [17]. This technique has been further developed in subsequent research papers, such as [3] and [8]. Overall, spectral techniques provide a different computational approach from spatial techniques and can be useful for training in the presence of a complex system graph.

K. Section Conclusion

The primary aim of this section is to suggest relevant computational techniques for integrating a proven computation module with an AI engine that is already in production. Modifying a preexisting implementation may be a more robust and supportable solution compared to building from scratch.

Spatial techniques, such as MARL, offer several preexisting implementations to choose from. However, finding a global optimum becomes more difficult as the number of agents representing processes grows larger, especially if the global objective function contains higher-order effects or piece-wise linear features. On the other hand, spectral techniques, based on the graph Fourier transform, can compress the learning solution to a more tractable size and tend to be more explainable based on the first few Fourier coefficients. Spatial techniques tend to come with "explainability" issues, as opposed to spectral techniques which are more "explainable" based on the first few Fourier coefficients. However, robust preexisting implementations may not be readily available for spectral techniques due to numerical issues when solving eigensystems for sparse matrices on large graphs. Certain computational heuristics may be required to work with graph Laplacian matrices containing a small number of extremely tiny negative eigenvalues due to the imprecisions in machine numerical representations. The net outcome is that, while spectral techniques appear to have certain computational advantages, using them to solve industry-sized problems may require customized implementations.

In practice, sacrificing certain computational gains and elegance from spectral techniques for a more practical initial implementation with fewer unexpected numerical issues may be worth the effort. Once the AI engine is trained, there is no requirement to continuously retrain it. Taking 10 or more times longer to train a GNN using spatial techniques may not be a showstopper, considering the overall gains from a much simpler integration of a preexisting implementation module. Eventually, the goal is to perform a spectral technique implementation as the problems of interest become bigger and more complicated.[5]

[4]The most recently updated list of relevant papers and implementation of GNN is available from https://github.com/thunlp/gnnpapers.

[5]The precise implementation details will be available from https://www.hedgespa.com.

V. Project Analysis

This section will discuss expected outcomes and risk and mitigation methods.

A. Expected Outcomes

Theoretical yield computation equations for a PV Panels to Converter to Electrolyzer system have been documented in detail by [14], which can be used to estimate the electrolyzer outputs based on various parameters such as the current, voltage, and operating temperature, as follows:

$$I = \frac{P_{EZ}}{U}$$

$$U = V_{cell} \cdot n_c$$

$$V_{cell} = \left(\frac{r_1 + r_2 T_{EZ}}{A_{EZ}} I\right) + s \log\left(\frac{t_1 + \frac{t_2}{T_{EZ}} + \frac{t_3}{T_{EZ}^2}}{A} I + 1\right) + 1.229$$

$$\dot{n}_{H_2} = \eta_F (n_c I / 2F)$$

$$\eta_F = \frac{(I/A_{EZ})^2}{f_1 + (I/A_{EZ})^2} f_2 (1)$$

These equations involve several parameters, including the electrolyzer current, voltage, Ohmic resistance parameters, electrode overvoltage coefficient, overvolt coefficients, number of cells, electrolyzer area, operating temperature, Faraday efficiency, and Faraday constant, as follows:

- I is the electrolyzer current;
- U is the electrolyzer voltage;
- r_1 is the Ohmic resistance paramter at $8.05 \times 10^{-5} \Omega m^2$;
- r_2 is the Ohmic resistance paramter at $-2.5 \times 10^{-7} \Omega m^2 / ^\circ C$;
- s is the electrode overvoltage coefficient at $0.185V$;
- t_1 is the overvolt coefficient at $-0.1002 m^2/A$;
- t_2 is the overvolt coefficient at $8.424 m^{2\circ} C/A$;
- t_3 is the overvolt coefficient at $247.3 m^{2\circ} C^2/A$;
- n_c is the number of cells in the electorlyzer, e.g. 6;
- A_{EZ} is the electrolyzer area, e.g. $0.25m^2$;
- T_{EZ} is the electrolyzer operating temperature, e.g. $80^\circ C$;
- η_F is the Faraday efficiency, which typical values range from 70% to 90% [13]; and
- F is the Faraday constant at $96,485 C/mol$.

The Faraday efficiency is a crucial parameter that depends on the specific chemistry of the system. Similar equations approximating theoretical production yields can be found in references such as [10] and [20], some of which may be based on modified assumptions on the equipment set. We can train a GNN with a simulation based on the chemical equations first before moving on to a physical system. This approach can have the added benefit of ensuring that the GNN is better equipped to handle rare out-of-sample “incidents” where even a physical system may not produce enough data to train the GNN effectively. It’s also worth noting that the closer the implied graph of the GNN resembles the chemical equations, the more likely the GNN will be able to produce sensible results.

An AI-driven control system is expected to push operating accuracy to well over 80% of theoretical limits, as discussed in [31]. Any consistent performance gaps should be explainable, such as material inputs that cannot be fed in continuously due to supply issues. The maximum yield should be defined as a limit that is achievable by the AI engine once all system frictions are removed. In other words, the maximum yield $\max_{\pi} \mathbb{E}_\pi[y(T;\pi)]$ from the previous section should be defined as a limit that is achievable by the AI engine once all solvabe system frictions such as supply chain issues of process input materials are removed.

The proposed project aims to demonstrate that the theoretical results can be achieved via physical experimentations. A demonstration unit will be built to provide evidence to the financiers, such as underwriters of asset-backed securities with tokenization of carbon credits, that the proposed technical setup will work in a physical plant, not purely based on simulations. The expected outcomes of the demonstration plant include verifying the theoretical yield, demonstrating the effectiveness of the AI-driven control system, and providing objective evidence for financiers. Its expected outcomes are summarized as follows:

1) Generator Size: According to one manufacturer’s specifications [28], a 5-megawatt commercial PEM electrolyzer can yield 83.75 kilograms of 99.999% high-purity hydrogen. Such equipment can be ordered based on a more customized specification. In other words, a 6-megawatt electrolyzer is expected to generate 100.5 kilograms of 99.999% high-purity hydrogen per hour.

In Asia today, solar mirrors and panels are regularly built on top of aquaculture ponds to generate electricity and lower evaporation. Any location receiving more than 2,000 sun-hours per year can be a potentially interesting target site.

While the generators may take 0.5 hour of warming up to reach operational efficiency, solar-thermal towers (with the right type of molten salt or thermal transfer oil) have the ability to continue operations for hours after sunset due to the very high temperature already accumulated in the heat circuit to generate steam. Therefore, it is conservative to assume that we have at least 12 hours of production per day, although some literature has suggested that the solar tower can continue to run for up to another 12 hours after sunset, allowing close to 24-hour production during the summer months [15].

As a result, the yield for a 6-megawatt PEM electrolyzer in one year is 328,599 kilograms or 328.6 metric tonnes of hydrogen if the plant can only run for 12 hours per day on average, but the yield can reach as much as 657,000 kilograms or 657 metric tonnes if the plant can run 24 hours per day. 329$\sim$657 metric tonnes per year is consistent with our production target of 1$\sim$2 metric tonnes per day of hydrogen. It should be noted that this is the maximum yield assuming all electricity produced is used to generate hydrogen. Depending on the pricing, some electricity may be sent to the power grid.

2) Water Intake: The amount of water required to produce 100.5 kilograms of hydrogen is 804 kilograms of pure water, which is roughly equivalent to 804 liters at ambient temperate.

Fig. 10. Solar Two Power Plant (National Renewable Energy Laboratory, Public domain, via Wikimedia Commons).

However, even if pure water is fed into the electrolyzer, there should be an allowance for a 10% loss due to production factors such as evaporation. So, we estimate 893 liters of pure water or roughly 15 liters per minute. This is comparable to the manufacturer's recommendation of 1,000 liters per hour, which typically includes an additional buffer to avoid commercial liabilities.

Do note that the flow rate of 15 liters of water per minute is similar to the intake flow rate of a large-capacity washing machine. In the worst case, even taking treated city water at such a flow rate is unlikely to draw significant environmental objections. Besides mixed water, we are exploring using either used industrial water or sewage after primary treatment, since a distillation step can be performed using the residual heat to extract the pure water. If we are allowed to extract 10% of pure water from any mixed water, we will require an intake of 10,000 liters of mixed water per hour to extract 1,000 liters of pure water per hour.

3) Heliostat Mirror Site Size: At the Solar Two Power Plant site near Barstow, California (as shown in Figure 10), a 10-megawatt solar plant requires a field size of 82,750 square meters. For 6 megawatts, we estimate the field size required to be 49,650 square meters.

Do note that Solar Two is using older technology from the 1980s. The efficiency of solar-thermal technology has since improved significantly so it is reasonable to treat this as an upper bound on the field size. Moreover, Bartow's latitude is 34.9 degrees north while most target sites are closer to the equator, which means that the solar radiation will be stronger in those locations due to their proximity to the equator. We expect to require fewer mirrors to achieve the same thermal output.

A field size of 49,650 square meters will require 2,758 solar mirrors in a standard 18-square-meter configuration. In a rectangular site configuration, that will take roughly 53×52 mirror units. Typical 18 square-meter mirrors come as a pair of 3 meters by 3 meters solar reflectors. Allowing 2 meters of margin for each mirror pair to minimize blockage from each other means that each mirror pair requires 5 meters by 8 meters of space, or roughly 110,240 square meters in total, translating to 332 meters on each side of a square configuration and 374.7 meters in diameter in a circular configuration. In other words, any regularly-shaped demonstration site of $1{\sim}2$ square kilometers for both hydrogen production and shipment consolidation will meet the commercial requirements, and still leaves plenty of room for safety margins and other environmental considerations such as settlement and cooling ponds if mixed water is used; typically, we are allowed to extract only $10{\sim}15\%$ of pure water and return the rest to the source "as close to ambient temperate as reasonable". Finally, NASA's Surface Meteorology and Solar Energy (SSE) 6.0 provides 1-degree latitude by 1-degree longitude grid-cell data over the globe, which is reasonably accurate as a starting point for the azimuth and elevation angle control software to calibrate the reflecting directions of mirrors.

4) Technical Goals of the Pilot Project: Our primary goal is to fine-tune the optimal production parameters based on physical experiments. For instance, is it better to maintain a higher temperature and aim for 24-hour production? Given that our goal is to implement smaller, self-running plants that may sit on different sites, there may be fewer maintenance issues (e.g. accelerated metal fatigue due to continuous high temperatures and pressure) by running the plants at lower temperatures with regular cooling down periods. While there is a large body of similar performance studies based on simulations in the scientific literature, there is no substitute for collecting the empirical physical data to fine-tune our performance models before replicating the plant design *en masse* at other locations.

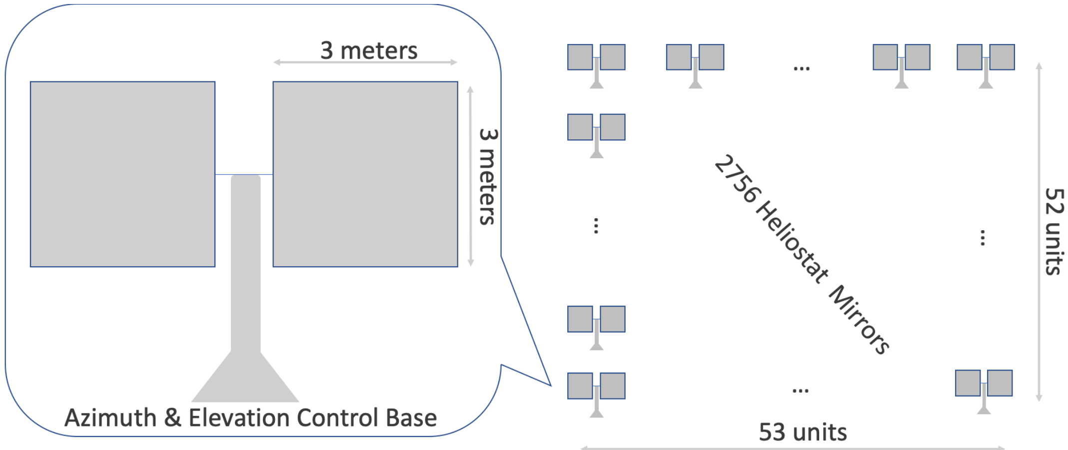

Fig. 11. Heliostat Schematic for Typical 3m×3m Solar Reflector Pairs.

5) Costs and Benefits: To meet all of our objectives above, our budget worksheet is estimating a total cost of ∼14 million US dollars over 3 years, primarily to pay for the equipment and construction, as well as local costs such as permits, water access rights, and environmental assessments. We target either producing up to 5∼10 megawatts of power or exporting 8∼13 metric tons of green ammonia/day or 1.4∼2.3 metric tons of green hydrogen/day, or a combination of both depending on the tariff rate offered by the utility company. This is expected to satisfy up to 2 kilograms/day of hydrogen demands per vehicle for a fleet of 1,000 hydrogen vehicles based on a highly achievable delivery price target and similar industrial uses. We will start with the blueprints based on standard and proven technologies with Technology Readiness Level (TRL) 7 or TRL8. The key innovation is how we put highly proven components together: for instance, by introducing the latest German high-temperature and high-pressure sensor technologies running with the HedgeSPA AI engine to drive the chemical processes to improve efficiencies and yield.

6) Work Packages: The specific collaboration with the Fraunhofer Institute will include the following R&D collaboration areas to address immediate needs for the initial demonstration facility. In most energy plants, sensors are usually placed at equal distances or in a topologically regular pattern to detect “fault signatures”. Our goal for the First Phase is to place the sensors outside the pipelines and reaction chambers. Our goal for the Second Phase is to place some of them inside the pipelines and reaction chambers. Each time we shall start with a lab-scale setup, primarily to avoid unnecessary changes and potential damage to the demonstration plant.

1) First Phase – Classic approach to set Project baseline:
   - Ready-made components;
   - Some microcontrollers outside reaction chambers (from AWS’ approved list);
   - Aiming for 4-pin connections with insulated connections;
   - Liquid/air cooling of microcontrollers and connections;
   - Ensure that the combination of sensors/actuators (pressure, temperature, viscosity, movements); and can meet the AI control and safety shutdown requirements.
2) Second Phase – R&D approach to improve efficiency and reliability:
   - More specialized, bespoke components;
   - Microcontroller embedded into the sensor package;
   - Insulate the entire unit with insulation materials;
   - Study temperature/failure characteristics as well as heat mitigation procedures; and
   - Where feasible, upgrade specific components in the overall design with the more advanced sensor packages.

This project is exciting to the Fraunhofer Institute in that we will be a good showcase to deploy IOT-style sensors and actuators controlled by our AI engine on AWS using Web Sockets or MQTT protocols. With a solid collaboration model, we may explore further collaboration with the Fraunhofer Institute on advanced high-temperature coating materials for solar-thermal applications and the use of their Graph Neural Network library to enhance the HedgeSPA AI engine.

7) Post-Demonstration Phase: After the demonstration phase, the financiers are interested in funding the full-scale production to meet up to 1 gigawatt in expected green energy demands that have expressed tentative interest and can begin delivery as soon as possible. As one example, we have identified a significant pool of institutional assets that will be interested in placing meaningful AUM into this project.

B. Risk and Mitigation

The technical processes involved are hazardous by nature. Considering this, our team members have identified the risks expected and possible risk management solutions as well as areas of minimum impact during operation to be:

1) Solar-thermal electrolysis – No known challenges, except for finding suitable sites with high solar radiation

that are not too far from a delivery hub and without unmanageable cyclone risks.

2) Similarly, while Polymer Electrolyte Membrane (PEM) is a proven technology, there may be issues gaining access to state-of-the-art polymer membranes currently manufactured by primarily Western companies for a plant targeting certain Asian buyers. Efficient production is the key to project success. Due to the current geopolitical environment, we must account for a non-zero probability that this project can be denied access to the most advanced electrode materials unless we agree not to deliver to potential customers in certain countries.
3) Alternative ways to transport/store hydrogen – We are looking at potentially using methylcyclohexane as the transfer agent, or using a carbon source to create synthetic natural gas. Successful delivery will depend on a solution to transport/store hydrogen at reasonable costs.
4) Green-bond Issuance – As mentioned, we do not have to issue in Hong Kong if the regulatory process is too complex. There is plenty of interest in such issuance in other jurisdictions, such as Australia, Singapore, or Japan.
5) Hyperledger Fabric Tokenization – No known challenges except for the need to deploy time/resources. Hyperledger Fabric is a public domain platform.

VI. Final Conclusions

The key drivers of success behind this proposed project are as follows:

1) Our team starts with a fully functional AI/GNN engine that has generated investment outperformance (positive returns and positive skewness) across 5 different markets or an outcome with less than 0.1% probability based on pure luck [12]. That means users of our attractive fintech platform can also be natural investors in our private investment projects in sustainable energy.
2) Scientific research linked to our AI algorithm has appeared at and was validated by top academic venues such as American Statistical Association and Association for Computing Machinery and will be published by the venerable scientific publisher Springer-Nature [12]. In particular, we are known for advanced research on predicting higher-order tail loss. These sophisticated risk management tools that were proven in managing energy assets should give additional confidence to future consumers who are interested in a total sustainability management solution.
3) The team has experience in energy investment projects at BlackRock and energy majors such as Chevron-Texaco. Several regulators have indicated that they are ready to approve digital assets related to sustainable energy faster than other asset classes, increasing the probability of project success.
4) We have several seasoned and proven investors in sustainable energy acting as our investors and advisors. The tokenization project was a finalist in the Tokenized Assets and Digital Securities (TADS) Award in 2020 and the Hydrogen Impact Investment Awards in 2022. These awards and name recognitions should build the confidence among future investors, without them it will be difficult to pull off sustainability projects on a large scale.
5) We work with identified demands of investors interested in our sustainability projects and demands of renewable energy customers keen to take delivery at or below certain achievable price targets, especially after government subsidies. In the current environment, investors are keen on sustainability projects, and consumers are required by their own regulator(s) to have a portion of their energy supply coming from renewables. These facts lend credibility to our green bonds which will be sweetened by a tokenization mechanism to attach high-integrity carbon offsets, increasing the chances of project success.

Einstein once said, “Insanity is doing the same thing over and over and expecting different results.” Global organizations have been driving green energy solutions for at least 2 decades. Adoption is often driven by former bureaucrats and project engineers from the energy majors, who are naturally biased toward doing things in the traditional way that they are comfortable with. We present a fresh if not disruptive approach that will hopefully drive adoption faster to reverse worrisome trends.

Instead of waiting for ground-breaking green technologies to emerge to make an impact, the most obvious bottlenecks seen today are integration and financing. We can start with proven components and put them together in innovative ways. Our engineers work side by side with financiers to work out the most innovative way to remove each and every roadblock with “divide and conquer”. There is now a once-in-a-lifetime opportunity to pull this off due to the dramatic change in attitudes among heavy users of fossil fuels, as a result of global geopolitical developments.

Finally, the carbon market in Asia failed due to a (mis)perception of widespread fraud. We can try the traditional way of fixing such a problem with heavy-handed regulations, which may kill off the nascent carbon market completely, or we can try to address this “trust issue” by using smarter technologies to create high-integrity carbon offsets. We have chosen the second approach and welcome other aspiring financial technology professionals to share any related and perhaps even better suggestions.

Acknowledgment

We thank the Fraunhofer Institute for Reliability and Microintegration or IZM (Dr. A. Ostmann, Dr. D. Manessis, Dr. M. Hempel), ILFA (Mr. T. Cramer), SignalGeneriX (Dr. T. Kounoudes) for the concept development of the electronic monitoring system, as well as Mr. E. Lee of Olympus Capital Asia for his insights on sustainable energy project finance.

REFERENCES

[1] W. Bernard Lee, Anthony G. Constantinides, and Tasos Kounoudes, "Green Hydrogen Production, Transportation, Securitization, and Tokenization", Presentation at *IARIA Green* 2022.

[2] Ahmet Aktaş, Yağmur Kirçiçek, "Solar System Characteristics, Advantages, and Disadvantages", Chapter 1, *Solar Hybrid Systems: Design and Application*, Academic Press, 2021, Pages 1-24, https://doi.org/10.1016/B978-0-323-88499-0.00001-X.

[3] Joaquin Bruna, Wojciech Zaremba, Arthur Szlam, and Yann LeCun, "Spectral networks and locally connected networks on graphs," in: *Proceedings of ICLR*, 2014.

[4] Nicos Christofides, "Bounds for the Travelling Salesman Problem", *Operations Research*, Vol. 20, No. 5 (Sep. - Oct., 1972), pp. 1044-1056 (13 pages). https://www.jstor.org/stable/169166?seq=1.

[5] Edward Farhi, Jeffrey Goldstone, and Sam Gutmann, "A Quantum Approximate Optimization Algorithm", MIT-CTP/4610, MIT Center for Theoretical Physics, 2014. https://arxiv.org/pdf/1411.4028.pdf.

[6] Thomas L. Gibson, and Nelson A. Kelly, "Predicting Efficiency of Solar Powered Hydrogen Generation Using Photovoltaic Electrolysis Devices", *International Journal of Hydrogen Energy*, Volume 35, Issue 3, February 2010, Pages 900-911. https://www.sciencedirect.com/science/article/abs/pii/S0360319909018448.

[7] Yongqing Fan, Zhan Zhu, Zhen Li, and Chenguang Yang, "Neural Adaptive with Impedance Learning Control for Uncertain Cooperative Multiple Robot Manipulators", *European Journal of Control (2023)*, doi: https://doi.org/10.1016/j.ejcon.2022.100769.

[8] Mathieu Henaff, Joaquin Bruna, and Yann LeCun, "Deep Convolutional Networks on Graph-Structured Data," arXiv preprint arXiv:1506.05163, 2015.

[9] Jing Huang, Jianyu Su, and Qing Chang, "Graph neural network and multi-agent reinforcement learning for machine-process-system integrated control to optimize production yield", *Journal of Manufacturing Systems*, Volume 64, 2022, Pages 8193.

[10] Hassan Ishaq, Owais Siddiqui, Ghaleb Chehade, and Ibrahim Dincer, "A solar and wind driven energy system for hydrogen and urea production with CO2 capturing," *International Journal of Hydrogen Energy*, 2021; 46: 4749-60. DOI: 10.1016/j.ijhydene.2020.01.208.

[11] Alina LaPotin, Kevin L. Schulte, Myles A. Steiner, Kyle Buznitsky, Colin C. Kelsall, Daniel J. Friedman, Eric J. Tervo, Ryan M. France, Michelle R. Young, Andrew Rohskopf, Shomik Verma, Evelyn N. Wang, and Asegun Henry, "Thermophotovoltaic efficiency of 40%", *Nature*, 13 April 2022.

[12] W. Bernard Lee and Anthony G. Constantinides, "Quantumized Graph Cuts in Portfolio Construction and Asset Selection", forthcoming in *Springer-Nature Transactions on Computational Science and Computational Intelligence*, 2022.

[13] N. Aaron Kelly, "Hydrogen Production by Water Electrolysis", Chapter 6 in *Advances in Hydrogen Production, Storage and Distribution*, 2014, Pages 159-185.

[14] Mohamed Nasser, Tamer F. Megahed, Shinichi Ookawaram, and Hamdy Hassan, "Techno-economic assessment of green hydrogen production using different configurations of wind turbines and PV panels", *Journal of Energy Systems*, 2022, 6 (4). DOI: 10.30521/jes.1132111.

[15] Salahuddin Qazi, "Solar Thermal Electricity and Solar Insolation", Chapter 7, *Standalone Photovoltaic (PV) Systems for Disaster Relief and Remote Areas*, Elsevier, 2017, https://www.sciencedirect.com/book/9780128030226/standalone-photovoltaic-pv-systems-for-disaster-relief-and-remote-areas.

[16] Fabian Scheepers, Markus Stählera, Andrea Stähler, Edward Rauls, Martin Müllera, Marcelo Carmo, and Werner Lehnert, "Temperature optimization for improving polymer electrolyte membrane-water electrolysis system efficiency", *Applied Energy* Volume 283, 1 February 2021, https://doi.org/10.1016/j.apenergy.2020.116270.

[17] David I. Shuman, Sunil K. Narang, Psacal Frossard, Antonio Ortega, and Pierre Vandergheynst, "The emerging field of signal processing on graphs: extending high-dimensional data analysis to networks and other irregular domains", *IEEE SPM 30 (2013)*, pages 83-98.

[18] Abishek Sriramulu, Nicolas Fourrier, and Christoph Bergmeir, "Adaptive dependency learning graph neural networks", *Information Sciences*, Volume 625, May 2023, Pages 700-714.

[19] Ljubiša Stanković, Danilo Mandic, Miloš Daković, Miloš Brajović, Bruno Scalzo, Shengxi Li and Anthony G. Constantinides, "Chapter 13: Portfolio Cuts", *Data Analytics on Graphs Part III: Machine Learning on Graphs, from Graph Topology to Applications*, Foundations and Trends® in Machine Learning: Vol. 13: No. 4, pp 499-511. http://dx.doi.org/10.1561/2200000078-3.

[20] Erik Van der Roest, Lennart Snip, Tak Fens, and Ad van Wijk, "Introducing Power-to-H3: Combining renewable electricity with heat, water and hydrogen production and storage in a neighbourhood." *Applied Energy* 2020; 257: 114024. DOI: 10.1016/j.apenergy.2019.114024.

[21] Fei Tan, Lili Zhou, Junwei Lu, Hongzheng Quan, Kaiyue Liu, "Adaptive quantitative control for finite time synchronization among multiplex switched nonlinear coupling complex networks", *European Journal of Control (2023)*, pages 23-37.

[22] Jie Zhou, Ganqu Cui, Shengding Hu, Zhengyan Zhang, Cheng Yang, Zhiyuan Liu, Lifeng Wang, Changcheng Li, Maosong Sun, "Graph neural networks: A review of methods and applications", *AI Open 1 (2020)*, Pages 57-81.

[23] "Average monthly hours of sunshine in Darwin (Northern Territory)", Weather and Climate, https://weather-and-climate.com/average-monthly-hours-Sunshine,Darwin,Australia.

[24] "Ashalim Power Station", Wikipedia, https://en.wikipedia.org/wiki/Ashalim_Power_Station.

[25] "Best Research-Cell Efficiency Chart", NREL, 2021. https://www.nrel.gov/pv/cell-efficiency.html.

[26] "Carbonomics: The Clean Hydrogen Revolution", Goldman Sachs Research, 07 Feb 2022. https://www.goldmansachs.com/insights/pages/carbonomics-the-clean-hydrogen-revolution.html.

[27] "Ivanpah Solar Power Facility", Wikipedia, https://en.wikipedia.org/wiki/Ivanpah_Solar_Power_Facility.

[28] "High purity 99.999% green energy powered ALK PEM wind solar hydrogen generator with electrolyzer", Alibaba, https://www.alibaba.com/product-detail/High-purity-99-999-green-energy_1600557711970.html.

[29] "PS20 solar power plant", Wikipedia, https://en.wikipedia.org/wiki/PS20_solar_power_plant.

[30] "SunShot Vision Study", US Department of Energy, 2022. https://www.energy.gov/eere/solar/sunshot-vision-study.

[31] "What is Accuracy", IGI Global, 2023. https://www.igi-global.com/dictionary/accuracy/329.

[32] "Wind Turbines Taller Than the Statue of Liberty Are Falling Over", Bloomberg, 23 January 2023. https://www.bloomberg.com/news/articles/2023-01-23/wind-turbine-collapses-punctuate-green-power-growing-pains.