Masteroppgave

NTNU
Norges teknisk-naturvitenskapelige universitet
Fakultet for informasjonsteknologi og elektroteknikk
Institutt for elkrafteknikk

Kristian Astad Dupont

Energy Certificate Trading using Distributed Ledger Technologies

Masteroppgave i Energi og Miljø
Veileder: Ümit Cali
Medveileder: Ugur Halden
Juni 2023

NTNU
Kunnskap for en bedre verden

Kristian Astad Dupont

Energy Certificate Trading using Distributed Ledger Technologies

Masteroppgave i Energi og Miljø
Veileder: Ümit Cali
Medveileder: Ugur Halden
Juni 2023

Norges teknisk-naturvitenskapelige universitet
Fakultet for informasjonsteknologi og elektroteknikk
Institutt for elkraftteknikk

NTNU
Kunnskap for en bedre verden

Abstract

English

As the world moves from a dependency on fossil fuels and gray energy, to a greener and more sustainable place, the importance of traceable energy increases. Within the European Union (EU), the European Energy Certificate System (EECS) has introduced Guarantees of Origin (GO)s to provide such energy provenance for consumers and producers. Currently, GOs are being bought and sold on a plethora of platforms, both peer to peer or through auctions. After the purchase, the concrete transfer of the GO is done on a separate system involving third parties needed to ensure the legitimacy of the trade. This thesis proposes that Distributed Ledger Technology (DLT)s may fulfill the role of both the marketplace and the trading mechanism. Specifically it looks at blockchain technologies, their functionality, use cases and flaws. Modern blockchains are designed for general information storage and trade, and the inclusion of smart contracts have enabled users to create complex conditions for how and when an asset should be transferred. Potentially, GOs could be both sold, bought and transferred on a blockchain, where the entire history of each certificate would be transparent for any auditing body to check.

Therefore, this thesis implements such a marketplace for energy certificates, using the Ethereum blockchain and with smart contracts written using Solidity. Furthermore, a co-simulation is created and presented in order to have realistic, but unique energy production and consumption data. This data is then converted to energy certificates, and one year worth of certificate production, consumption and trading is done through the smart contract. Through the use of the web3py Python library, calls to the smart contract can be automated and we are able to simulate a complex network of fifteen nodes stationed in Germany, Spain and England. The marketplace implementation successfully traded blockchain based certificates, but it still lacks some features for an accurate real world representation. Mainly being run on an extensive, but private blockchain, and increasing the number of network nodes.

Norsk

Ettersom verden endrer seg fra å være avhengig av fossile brennstoff og fossil energi, til å være grønnere og mer bærekraftig, øker viktigheten av å kjenne energiens opphav. EU har introdusert det Europeiske Energi Sertifikat system (EECS), og gjennom det opprinnelsesgarantier. Per nå blir opprinnelsesgarantier kjøpt og solgt på en rekke forskjellige plattformer, både gjennom en-til-en handler, og gjennom auksjoner. Etter handelen foregår den konkrete overførselen av opprinnelsesgarantiene på et separat system som involverer tredjeparter for å forsikre seg om legitimiteten av handelen. Denne avhandlingen foreslår at "Distributed Ledger Technologies" (DLTs) kan ta på seg rollen som både markedet og overførings-mekanismen. Helt konkret sed den på blokkjedeteknologier, deres funksjonaliteter, bruksområder og feil. Moderne blokkjeder er designet for å kunne både lagre og flytte generell informasjon, og etter inkluderingen av "smart contracts" er det blitt mulig for brukeren å sette komplekse krav når og hvordan en ressurs skal flyttes i blokkjeden. Potensielt kan opprinnelsesgarantier bli solgt, kjøpt og forlyttet på blokkjeden, hvor historien til hver enkel garanti kan ville vært synlig for alle som ønsker å sjekke dens legitimitet.

Derfor implementerer denne avhandlingen en slik markedsplass for opprinnelsesgarantier. Dette gjøres ved hjelp av Ethereum blokkjeden og "smart contracts" skrevet i Solidity. I tillegg lages en ekstra simulering for å ha realistiske, men unike energiproduksjons- og konsumerings-data. Denne dataen blir representert med opprinnelsesgarantier, og ett år med garantiproduksjon, konsumering og handel utføres gjennom den smarte kontrakten. Ved så å bruke Python biblioteket Web3Py kan man automatisere funksjonskall til den smarte kontrakt, og vi kan simulere et nettverk bestående av femten noder fra Norge, Spania og Tyskland. Det implementerte markedet handlet suksessfullt med garantier, men mangler fortsatt noen egenskaper for å kunne representere realistisk handel på en god og konsekvent måte. Hovedsakelig mangler den å bli utplassert og testet på en utbrett, men privat blokkjede, og å operere med flere simulerte noder.

Contents

List of Figures

List of Tables

Acronyms

ABI Application Binary Interface. 38, 39

AIB Association of Issuing Bodies. 11–13, 35, 36, 58

BRP Balancing Responsible Party. 8

CAPEX Capital Expendature. 6

CDP Carbon disclosure Project. 11

DER Distributed Energy Resources. 3, 8

DLT Distributed Ledger Technology. i, 1, 4, 7, 14, 20, 35, 59

DSO Distributed Systems Operator. 8

EEA European Economic Area. 12

EECS European Energy Certificate System. i, 11, 12, 58

EU European Union. i, 3, 11, 12, 58

EV Electric Vehicle. 8

EVM Ethereum Virtual Machine. 23, 24, 37, 38

GHGP GreenHouse Gas Protocol. 11

GO Guarantees of Origin. i, iv, 9–13, 26, 27, 35, 36, 58

ICT Information Communication Technology. 4

IREC International Renewable Energy Certificate. 9, 11

IRENA International Renewable Energy Agency. 6

LCOE Levelized Cost of Energy. 6

LEM Local Energy Market. 7, 8

LGC Large-scale Generation Certificates. 11

MERRA Modeling-Era Retrospective Analysis for Research and Applications. 27, 29–32

NFT Non-Fungible Tokens. 18, 36, 37, 45, 46, 48, 57, 59

OPEX Operational Expendature. 6

pBFT proof of Byzantine Fault Tolerance. 16, 19–21

PoA Proof of Authority. 16, 17, 19, 59

PoI Proof of Identity. 16, 17, 59

PoS Proof of Stake. 14, 17–22, 24, 25, 59

PoW Proof of Work. 14, 18–22, 24, 25, 39, 59

PoX Proof-of-X. 18–20

REC Renewable Energy Certificate. 9–11

REGO Renewable Energy Guarantee of Origin. 9, 12

RES Renewable Energy Resources. 4

ROC Renewable Obligation Certificates. 9

SFT Semi-Fungible Token. 18, 36, 37, 48, 57, 59

SGAM Smart Grid Architecture Model. 8

1 Thesis motivation

The energy sector is undertaking radical changes, motivated by the urgency of climate action to mitigate the climate crisis. There is a global aim to decrease the worlds dependency on fossil fuels, mitigate emissions and improve emission accountability. [1] [2] To improve the accountability, and promote green energy production, several nations and unions of nations have introduced energy certificates in order to distinguish green energy from gray energy.

Simultaneously, there is rapid technological development on the DLT front, with large actors such as Ethereum working rapidly to improve the security and scalability of their systems. [3] While this technology is no longer brand new, it is still not majorly adopted and implemented, and the landscape for innovation and new use cases is flourishing. DLTs are already widely explored in the context of local energy markets and peer-to-peer trading, but less so in the context of certificate trading.

Specifically, this thesis wishes to look at certificate trading at a macro scale, working with a number of certificates that requires the solution to be scalable. Overall, the goals of this thesis is to:

- Providing the background information required to comprehend the necessity for, and the status of energy certificates, and how they may be represented on a blockchain.
- Evaluate token standards for the use case of an energy certificate representation on a blockchain.
- Develop and test a marketplace proof of concept for energy certificates, in which certificate information is traded, transferred, and stored on the blockchain.
- Propose and implement a method for the production of certificate time series data by developing a co-simulation for the certificate market.

If this thesis succeeds in fulfilling all the abovementioned goals, we believe that it will provide the academic community with:

- A comparative analysis of token standards when used to represent energy certificates on the blockchain.
- Simulation results of the production, consumption and trade of energy certificates based on the blockchain for select parts of Europe.
- An improved and expanded overview of the intersection between the energy sector and the blockchain space. Specifically with regards to energy certificates.

The background information found in the following sections and subsections under section 2 and section 3 was first gathered for the specialization project leading up to this thesis. The information has been updated, paraphrased and rewritten to fit this thesis.[4]

2 Introduction

Currently, the climate crisis is a widely global and imminent problem, and energy sector is facing a critical need for transformative change. The production of electricity and heat stands out as the leading contributor to CO2eq emissions [5], necessitating an imminent shift from fossil fuel-based production to environmentally sustainable alternatives. Earlier this yeat, the UN Sharm el-Sheikh Climate Change Conference was held, and it was accompanied by a dire warning from the UN Secretary-General about the potential for a "climate hell".[6] Effective and actionable energy policies must be implemented by central governments, while industrial compliance and innovation are equally vital to facilitate a swift and seamless transition within the energy sector. As a prominent platform for global policymaking, the United Nations has outlined a set of future objectives to foster a sustainable world. These objectives encompass various aspects, including "Affordable and Clean Energy" (Goal 7), "Industry, Innovation and Infrastructure" (Goal 9), "Sustainable Cities and Communities" (Goal 11), and "Climate Action" (Goal 13) [7].

2.1 The Digital and Green shift

The Green Shift signifies a crucial transition from fossil fuel dependencies to renewable energy sources. It represents one among numerous steps towards achieving the UN's sustainability goals and paving the way for a livable future. It is a multifaceted process, involving both a plethora of sectors; housing, energy, transportation, agriculture, industry, and both governmental and private actors. The green shift has become an integral part of good large scale management, with governments proposing policies aimed at accelerating the transition, and companies using the green shift in their marketing. [2, 8]

With the green shift, the energy sector is likely to change drastically. The figure in Figure 1 is included to show the scale of this impact, and how integral an energy transition is to the green shift. The figure illustrates the global emissions broken down for the year 2016, categorized by their respective sources. When considering the emissions from transport as a part of the energy emissions, the energy sector accounted for approximately 73% of the world's total emissions. Within the energy sector, a significant portion of these emissions, around 40 out of the 73 percentage points, can be attributed to the utilization of electricity and heat for residential and industrial purposes. This emphasizes the crucial role that transitioning to sustainable energy alternatives can play in reducing overall emissions and mitigating the impact of climate change.

With the advancement of digital technologies, IoT devices, and information and communication technologies, a new concept called the "Digital Green Shift" has emerged alongside the existing Green Shift [9]. The unpredictable nature of renewable sources like wind and solar power has highlighted the need for innovative approaches in load distribution, energy flexibility, and storage to address the fluctuations associated with their increased utilization. Digital solutions and communication technologies play a vital role in optimizing the coordination of consumption

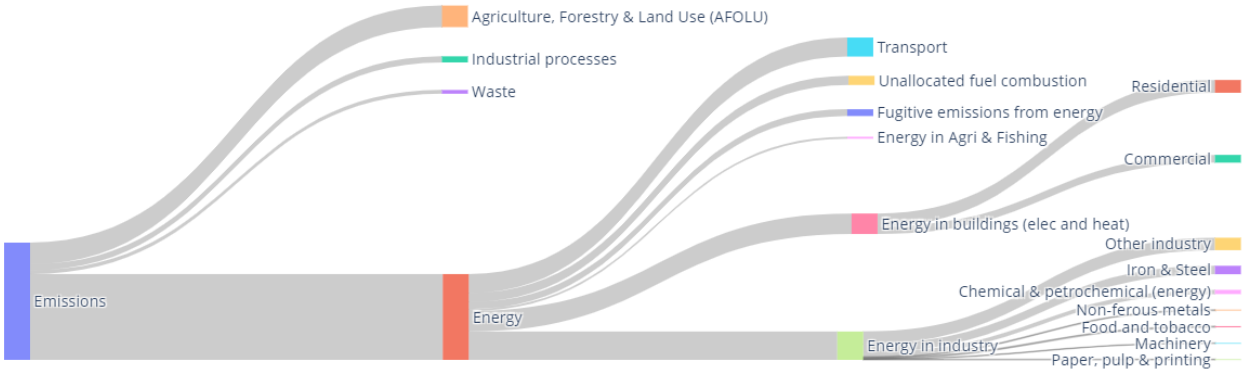

Figure 1: Global emissions by source in 2016. Data from [5].

and production. The energy sector is known for generating vast amounts of data, making it a prime candidate for the implementation of artificial intelligence (AI) and machine learning (ML). AI has been explored in various applications, including photovoltaic panels, hydrogen production, and energy planning and control. [10] By leveraging AI and ML technologies, these processes can be enhanced to improve the efficiency and effectiveness of renewable energy utilization, and improve the efficiency and speed of the Green shift.

2.2 The five Ds of the power system.

Between the late 1970s and the late 1990s, the political landascape surrounding the energy market changed significantly. Many Western countries experienced a wave of privatization and liberalization, resulting in the conversion of previously state-owned services into tradable and commercial commodities. The energy and electricity sectors were not excluded, and both individual countries and international unions started deregulating their power markets. In the United States, the deregulation process began with President Jimmy Carter pushing through two energy deregulation packages, one in 1978 and another in 1980 [11]. Great Britain followed suit with the Electricity Act of 1989 under the leadership of Margaret Thatcher [12], while the EU initiated the liberalization of their energy markets in 1996 through a directive focused on internal electricity markets. [13] Since the liberalization of power markets, the energy system has progressively moved towards inclusive power markets, where power production and trading are no longer restricted to governmental actors. The liberalization and deregulation of the power market was the first of what has since been called the five Ds of power systems/electrical systems [14]. The five Ds represent distinct stages in the development of power systems, and can be observed in Figure 2. Although the development of each "D" is continuous, and the the time frames vary, they can be thought to follow a roughly chronological order starting with Deregulation, and ending with Democratization.

Decentralization and decarbonization are integral components of the green shift. Decentralization refers to the increasing utilization of Distributed Energy Resources (DER)s by smaller companies and even individual households. This trend is fueled by the availability of cheaper small-scale power sources, particularly photovoltaic panels, enabling energy production to shift away from centralized producers and empowering consumers to generate their own energy [15]. The rise in electric vehicle sales further contributes to in-house energy storage capabilities [16],

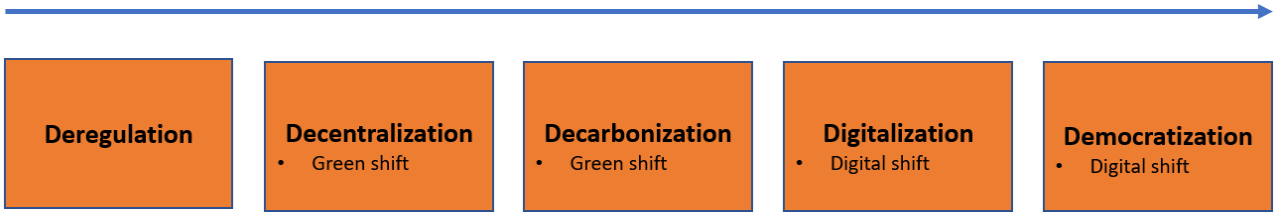

Figure 2: The five Ds of power systems.

and the ability of prosumers to produce energy not only for personal use but, but also to supply it to the grid is changing the fundamental architecture of the power grid.

Decarbonization involves the integration and development of green and Renewable Energy Resources (RES). Although the majority of global energy consumption continues to rely on unsustainable fossil fuels, there has been a global commitment, exemplified by the Paris Agreement [1], to prioritize carbon-neutral energy sources and reduce dependency on fossil fuels. Figure 3 illustrates the cumulative and yearly global energy consumption by source. The graph on the right reveals the predominant reliance on coal, gas, and oil. However, since the 2010s, there has been exponential growth in solar and wind energy consumption, while the consumption of oil and coal appears to have plateaued or even declined.

Digitalization and democratization are belong to the digital shift of the energy market. These changes are currently being researched and primarily implemented through pilot projects. Digitalization of the power market involves the integration of DLTs, smart meters, Information Communication Technology (ICT)s, machine learning, and artificial intelligence. A digitalized market serves as a foundation for smart automatic trading, optimized consumer strategies, and optimized social welfare. Parts of the digitalization of power systems are already in use, such as employing advanced AI and optimization techniques to forecast demand in populated areas [14]. Digitalization also plays a crucial role in the development of smart cities and local energy markets, enabling internal trading and storage to accommodate the increased penetration of RES [17].

Democratization focuses on expanding user participation. It builds upon the concept of decentralized production, expanding the liberties of the free market to local prosumers that can generate and trade energy amongst themselves. Through advancements in DLTs and the implementation of digital energy markets, prosumers and consumers gain direct influence over the price and value of energy, challenging the monopolistic nature of the current market. Democratization encourages prosumers to take an active role in their energy consumption and production, as well as consider the social and environmental impact of their energy use [18]. Increased consumer participation has also been found to mitigate the "not in my backyard" sentiment often associated with renewable energy projects. [19]

2.3 The transition to renewable energy

Over the past decade, global investments in renewable energy have shown consistent and rapid growth. The installed capacity for renewable power production increased from 1.44 GW in 2012 to 3.06 GW in 2021, resulting in more than a doubling of the capacity over a ten-year period. [21]

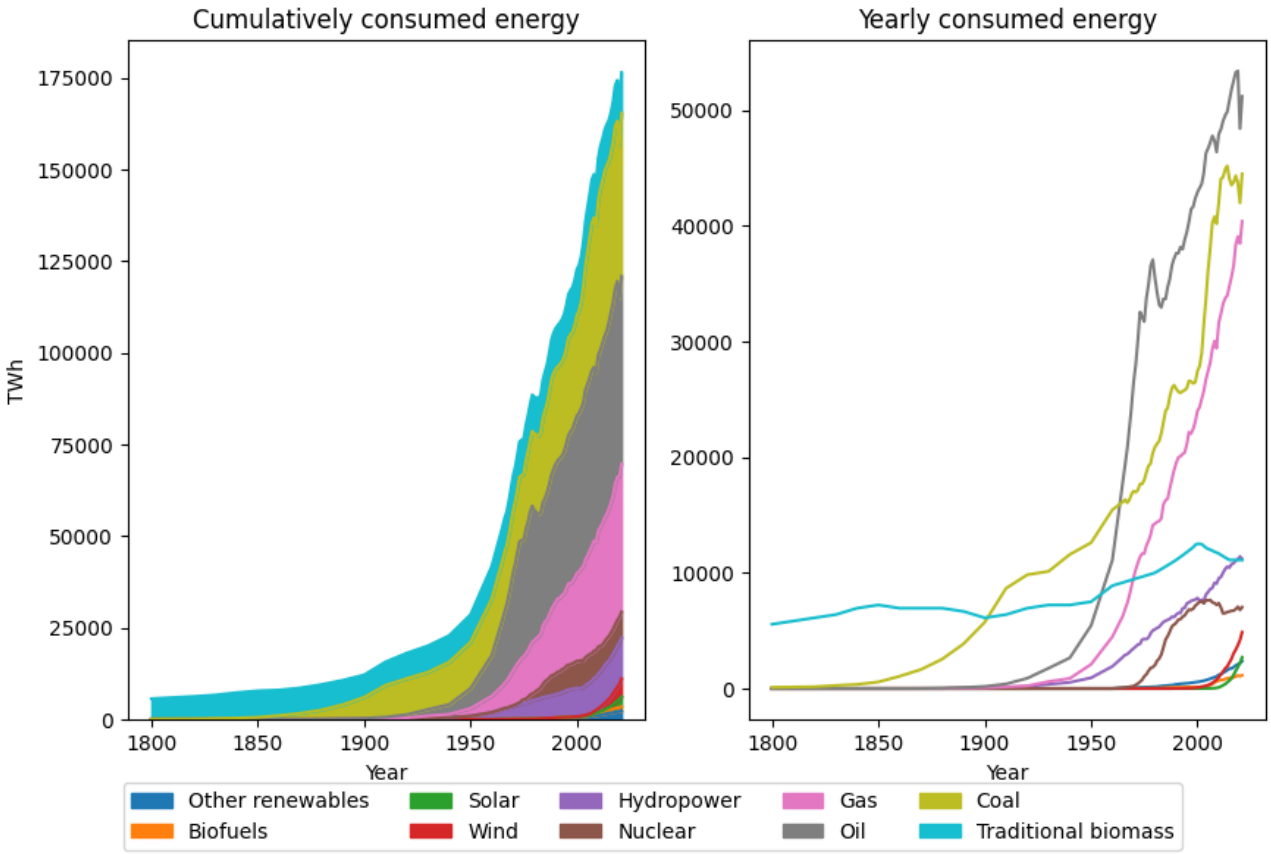

Figure 3: Cumulative and yearly consumed energy by source. Data from [20].

On average, the installed renewable capacity grew by 8.7% during this timeframe, demonstrating a steady exponential increase. In contrast, newly installed fossil fuel capacities have been declining rapidly since around 2015. From 2020 to 2021, newly installed fossil energy capacities were insignificant compared to newly installed renewable capacities [22]. While hydropower remains the largest source of renewable energy production, solar technologies dominate the market for newly installed capacities [23]. In 2018 and 2019, solar technologies accounted for as much as 56.6% of newly installed renewable technology capacities. The development of yearly installed capacity can be observed in Figure 4

The renewable energy sector relies on various sources including hydropower, wind power, solar power, biomass, and other sources such as tidal and geothermal power. In 2021, the distribution of renewable energy production among these sources was as follows: hydropower accounted for approximately 54%, wind power for 23%, solar power for 13%, biomass for 8%, and other sources contributed around 2% [24]. It's worth noting that the share of hydropower in renewable energy production has decreased compared to previous years. The figure in Figure 5 shows the development of renewable energy production over the last 41 years. The graph illustrates significant increases in energy generation from solar and wind sources, reflecting the substantial investments made in these sectors in recent years.

The significant increase in renewable investments can likely be credited primarily to the substantial reduction in renewable technology costs rather than purely environmental considerations.

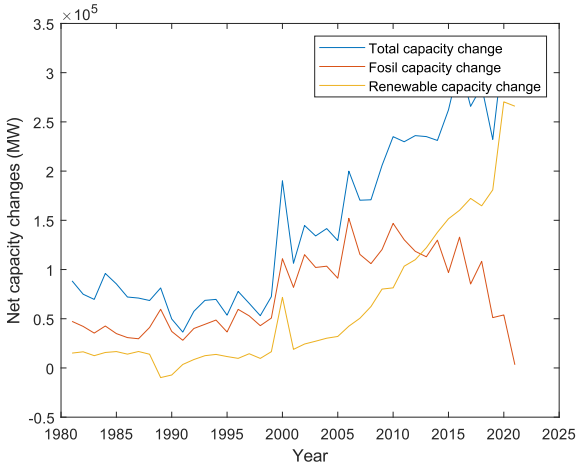

(a) Yearly installed energy capacity. [22]

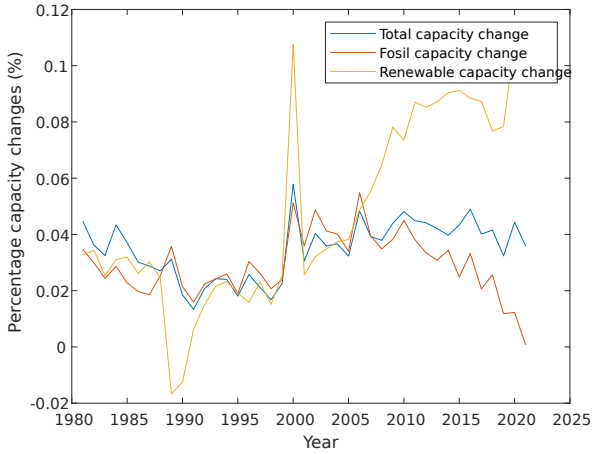

(b) Yearly installed energy capacity by percentage of existing capacity. [22]

Figure 4: Globaly installed energy capacities.

Still, the investment boom has likely has a significant impact on the long-term trajectory of the energy sector. One commonly used metric for evaluating and comparing the profitability of different energy sources is the Levelized Cost of Energy (LCOE). The LCOE represents the average cost of producing one unit of energy over the lifetime of a particular technology, encompassing both the Capital Expendature (CAPEX) and Operational Expendature (OPEX) of a system.

Over the past decade, the LCOE of most renewable sources, particularly solar and wind, has experienced a significant decline. According to the International Renewable Energy Agency (IRENA) report on global LCOE trends from 2010 to 2021, conventional solar technologies have seen their costs decrease to approximately 11% of their 2010 levels when considering the entire lifespan of photovoltaic panels.[25] Similarly, concentrated solar power costs have dropped by around 68%, but is now out competed by conventional solar power. Conventional wind power costs have decreased by 68%, and offshore wind power costs have decreased by 60%. It's important to note that hydropower and geothermal power are outliers in this trend, as their prices have increased over the same period. However, they were characterized by low initial costs, and the increase in LCOE is negligible compared to inflation.

While the LCOE of specific projects can vary significantly, the provided averages indicate that solar, wind, and hydro technologies can now compete with or even be cheaper than traditional fossil fuels when considering their entire lifespan [25]. It's worth mentioning that LCOE calculations often rely on various assumptions, particularly when estimating LCOE values for projects in 2021. Nevertheless, other sources confirm these numbers, showing similar or slightly higher LCOEs for renewable sources [26, 27].

In addition to cost-related challenges, the energy transition faces inherent public dissonance. While the concept of a greener shift is generally supported, it encounters resistance when it comes to specific and concrete initiatives. In a review conducted by [28], they analyze published surveys and experiments to provide an overview of public opinion on renewable energy. The

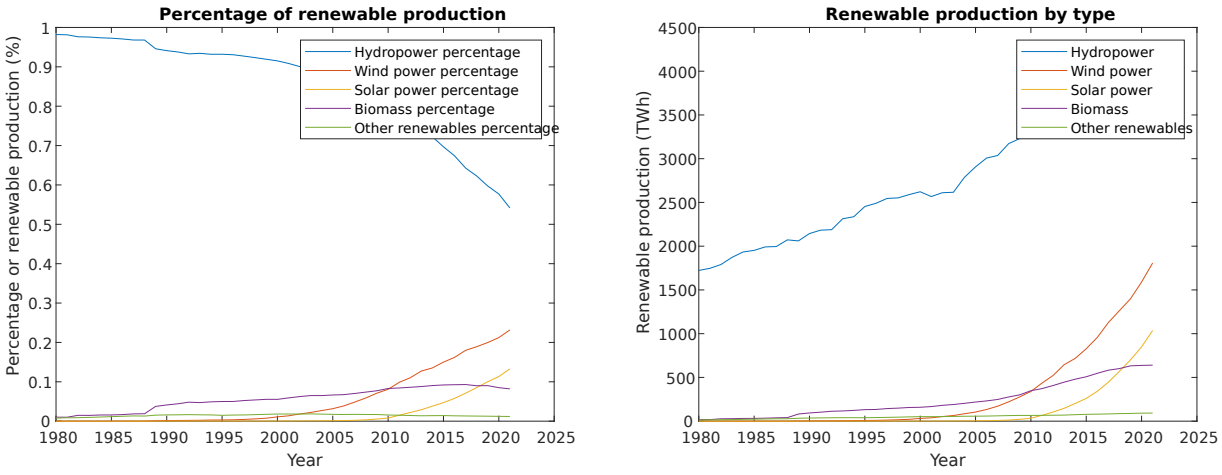

(a) Renewable energy source as a percentage of the total renewable energy production. [24]

(b) Renewable energy production by type. [24]

Figure 5: Renewable energy production by source.

review encompasses 19 surveys conducted in geographically distributed locations, addressing various aspects of renewable energy implementation. The findings reveal a mixed response from the public. Among the 19 surveys analyzed, eight reported positive responses, indicating support for renewable technologies. Nine surveys received partially supportive responses, and two surveys reported negative responses towards renewable energy initiatives. The paper presents a comprehensive table which identifies common barriers to the acceptance of renewable technologies. Some frequently encountered obstacles include a lack of knowledge about renewable technologies, personal interests related to affected resources or policies, concerns about noise, and demographic factors.

2.4 Recent changes within the power market

The traditional power market has been characterized by the trade, transmission, and distribution of electricity from generation to consumption. Typically, it consists of five main actors: the energy generator, the transmitter, the distributor, the retailer, and the consumer. However, the power market is undergoing significant changes due to the introduction of power electronics and the rise of prosumers. In the conventional top-down market structure, power primarily flows from large-scale power plants or centralized wind farms, however, there is a growing trend towards a more multi-directional flow of power in the future. Grid infrastructure models, as discussed in [29], can be classified into three categories: "bulk," "micro," or a combination of both. Bulk models resemble transmission systems that handle power delivery and distribution on a national scale, while micro models focus on smaller, self-contained distribution networks that act as separate clusters interconnected mainly by the larger bulk network.

The micro model is well suited to describe the contemporary idea of Local Energy Market (LEM)s. LEMs have gained significant attention in recent years, as evidenced by the growing number of impactful papers focused on this topic.[30] These markets are often explored together with the use of DLT applications, as the decentral nature of LEMs correspond well with the decentral nature of DLTs. A local energy market can be defined as a platform where residential

customers within a close-knit geographic and social community can trade locally generated renewable energy.[31] While local energy markets are typically connected to the larger energy grid, they offer the advantage of reducing reliance on central energy production.

The transition towards LEMs, and the integration of DERs are may be crucial for achieving a sustainable energy future. However, recent literature has highlighted several challenges associated with this transition. In a literature study conducted in [32], five main challenges are identified. First, the increased intermittency of renewable energy sources requires a higher capacity for balancing power supply and demand. Second, the rapid changes in voltage levels associated with DER integration can lead to overvoltages in the grid. Third, the reverse power flows from DERs can result in unpredictably high short-circuit currents. Fourth, detecting faults and ensuring protection in a system with high DER penetration can be challenging. Lastly, unpredicted islanding mode operations.

As a way to meet these challenges, the study explores the potential for increased flexibility and flexibility services. Flexibility in the context of the power grid refers to the ability of the grid to adapt to changes in power supply and demand. Flexibility services are designed to provide this adaptability, and are characterized by factors such as the power magnitude and rate of change they can offer, as well as their duration, energetic capacity, location, and response time. [32]

The literature review proposes three subcategories for flexibility: stationary storage and flexible loads, mobile storage, and flexibility contracts and markets. Stationary storage can for example involve the installation of large battery packs by the Distributed Systems Operator (DSO) or LEM to smooth out power fluctuations and avoid grid upgrades. Mobile storage may refer to Electric Vehicle (EV)s and their potential to provide flexibility through smart charging strategies, such as coordinated charging or vehicle-to-grid capabilities. Such smart charging strategies can enhance grid flexibility by mitigating power spikes caused by uncoordinated charging. Lastly, flexibility contracts and markets encompass a range of services provided by and to consumers and prosumers by the DSO or Balancing Responsible Party (BRP). This includes initiatives like demand response programs consensual load curtailment.

As the power market is undergoes the digital green shift, literature has started to view the power market as a more multi-purpose market, which no longer is limited to only energy trades. The paper presented in [33] proposes a three-layered market segmentation framework comprising power transactions, data transactions, and financial transactions. Power transactions refer to the physical exchange of electricity on the grid, while data transactions involve the sharing and trading of various data related to energy production, storage, labeling and consumption. Financial transactions encompass the exchange of financial instruments and contracts related to energy transactions. The paper focuses on the cybersecurity aspects of blockchain-based energy markets, envisioning that data and financial transactions can be securely conducted on the blockchain, while energy transactions are conducted on the physical grid. The figure in Figure 6 provides an overview of current and potential market services and their respective segments within this framework.

An alternative, yet similarly structured, three-layered segmentation can be proposed as: (1) a physical layer, (2) a certificate layer, and (3) a transactional layer. In this model, the physical layer encompasses the concrete input and withdrawal of energy, along with the associated grid infrastructure. This layer is similar to the physical layer defined in the Smart Grid Architec-

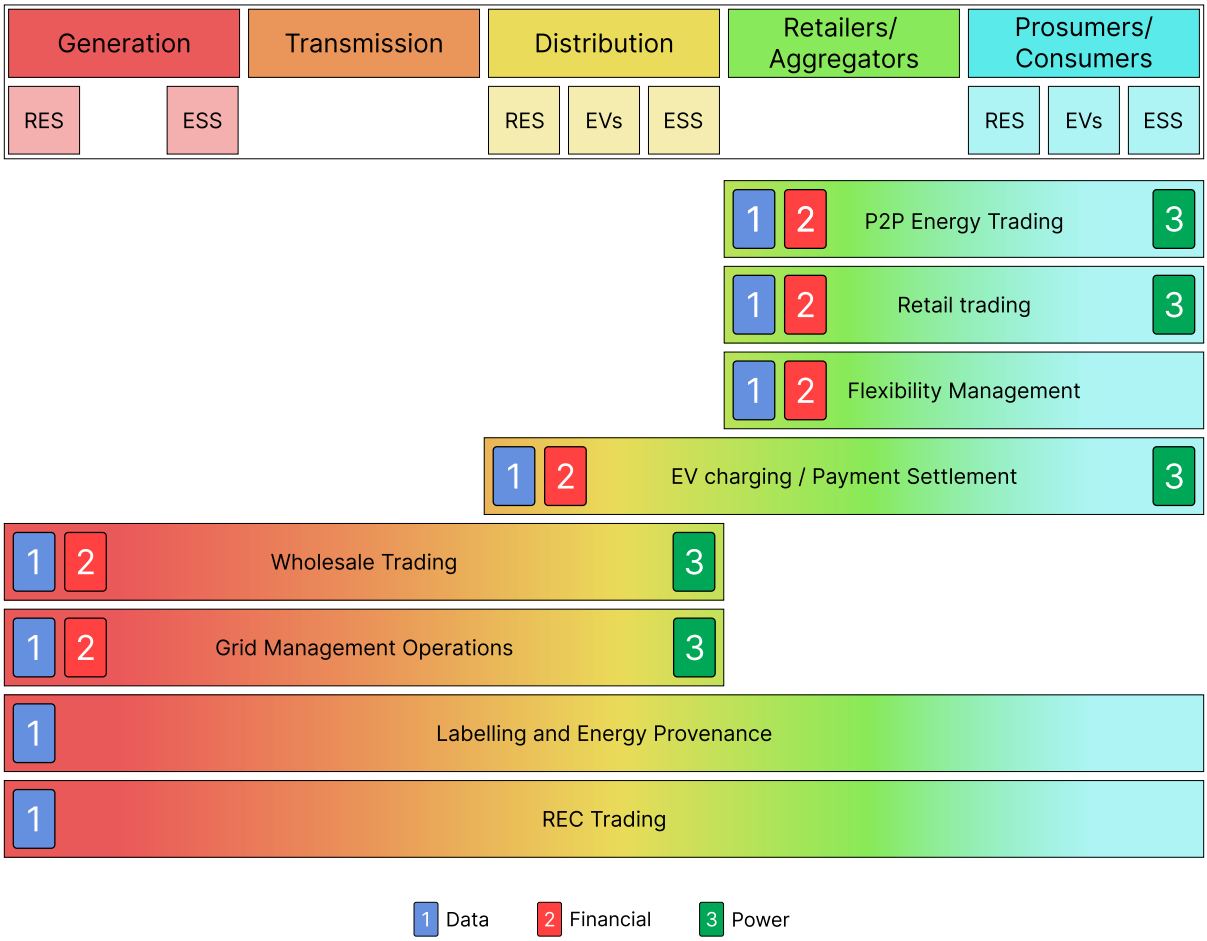

Figure 6: The current power market, with current and potential market services, divided by market location and market segment. The figure is directly adapted from [33].

ture Model (SGAM) [34]. The certificate layer represents the digital realm where energy is represented and the trade and tracking of energy provenance occur. It serves as a digital layer for trading energy certificates such as Renewable Energy Certificates (RECs) and Guarantees of Origin (GOs). More details on RECs and GOs are provided in subsection 2.6. Finally, the transactional layer encompasses financial transactions, potentially utilizing the blockchain as the financial ledger.

2.5 Renewable Energy Certificates/Energy Attribute Certificates

Distinguishing between electricity produced by renewable sources and fossil fuels becomes impossible when electricity moves from production to transition. Similarly, once processed biogas enters mixed transit, it becomes indistinguishable from fossil natural gas. To address this issue, governmental institutions have implemented certificates to assist clean energy producers in documenting their clean energy production and enabling consumers to track their clean energy consumption. These certificates have been globally adopted and include European GOs, US Renewable Energy Certificate (REC)s, International Renewable Energy Certificate (IREC)s, British Renewable Energy Guarantee of Origin (REGO)s and British Renewable Obligation

Certificates (ROC)s. RECs are entirely digital assets, which makes them well-suited for trade and export. They are not subject to the physical and geographical limitations associated with energy trade, earning them the title of "the currency of renewable energy markets". [35] RECs provide an opportunity for customers and consumers located far from renewable sources to support and promote the production of green energy over fossil fuel-based energy, effectively distinguishing the value between these two energy sources.

The use of different RECs to facilitate information tracking and promote renewable energy production is expected to grow nearly tenfold between 2020 and 2030, as the market is projected to increase from $10.2 billion to $100.96 billion [36]. Currently, the North American market holds the largest share, but the European market is the fastest-growing, with the most optimistic forecasts. The fast growth indicates the important role that RECs play in funding and incentivizing the development of green energy, and the green shift. However, there are concerns regarding the potential for RECs to enable greenwashing, which may hinder sustainable development and improved customer practices. [37] Greenwashing refers to the practice of ""misleading consumers about the environmental performance of a company as a whole"[38]. In other words, greenwashing is the act of falsely portraying ones environmental impacts as smaller than it really is. [39] It is challenging to categorically associate RECs with other forms of greenwashing, since they do provide direct incentives for energy companies to prioritize greener energy sources. Still, arguments have been made that RECs allow companies to present their energy usage as greener than it actually is. It's important to recognize that purchasing RECs does not immediately clean the energy mix at the moment of consumption; instead, it allocates emissions to parties that may not be able to afford to, or are less incentivized to operate sustainably [37]. Critics have scrutinized RECs for providing a low cost option for perceived emission reductions, relieving large scale emitters of the responsibility of actually improving their sustainability. [40]. Additionally, some argue that RECs can assist Bitcoin mining to appear more sustainable, even though Bitcoin mining inherently is purely consumal in nature, and does not currently return any tangible benefit to the society. [41]

The origin of RECs can be traced back to 1996 when the state of California initiated discussions regarding a state-wide Renewable Portfolio Standard. Following, in 1998, the first green power market was established, enabling Californians to engage in the trading of energy exclusively derived from renewable sources. In 1999, "green tickets" were introduced into California's gray power market. This involved bundling wholesale energy with proof of renewable sources, thus providing the retail market with access to dedicated green energy. [35]

The North American REC market is, as mentioned, the largest in terms of monetary value. In the United States, these certificates are commonly referred to as just RECs, or sometimes US RECs, although they are also widely used in Canada. The US market consists of two distinct REC markets: the compliance market and the voluntary market. The compliance market caters to customers who are obligated by policy regulations to meet specific Renewable Portfolio Standards, while the voluntary market is open to anyone who wishes to participate and purchase RECs to reduce their CO2 emissions voluntarily. [42] The voluntary market is secondary to the compliance market, and deals in excess certificates not sold on the compliance market. A notable difference between US RECs and for example GOs is that US RECs do not specify the specific source of energy represented by the certificate. Instead, they simply indicate that the energy meets renewable standards according to US regulations. However, there are

some state-specific RECs that focus exclusively on solar energy. These solar-specific RECs are referred to as solar RECs or sRECs [43].

In addition to the United States, several different types and standards of RECs have emerged globally. Among them, two international standards and one national standard complement the North American standard. The largest international standard is the IREC, which is utilized in countries accounting for over 51% of the world's energy consumption. IRECs are primarily used in Eastern Asia, Russia, and South America, with increasing adoption in African countries [44]. The International REC Standard Foundation, a non-profit organization, oversees the IREC standard. It ensures that consumers of IRECs can be confident that their purchased certificates are recognized by major sustainability standards, including the GreenHouse Gas Protocol (GHGP) and the Carbon disclosure Project (CDP) [45]. Other large scale certificate standards are the Aurstralian National standard, called Large-scale Generation Certificates (LGC)s, and the European GO.

2.6 Guarantees of Origin and the European Energy Certificate System

The EU has created the EN16325+A1 standard for GOs, which serves as a certificate to label energy based on its source. GOs are tradeable assets, with each certificate representing one MWh of energy. While the EU standard governs GOs, they are adopted not only by all countries within the European Economic Area but also by Switzerland. Currently, GOs can be issued in 28 countries, but are also purchased by some non-issuing countries. Each issuing country is considered a separate domain, except for Belgium, which consists of four domains. Each domain has at least one qualified issuing body responsible for issuing GOs. Examples of these issuing bodies include Statnett in Norway, the Umweltbundesamt (the German Environmental Agency) in Germany, and the European Energy Exchange in France. These issuing bodies operate under their respective domain protocols, which, although sharing similarities, have distinct characteristics. Each domain protocol includes an overview of the legal and regulatory framework within the domain, with for example protocols for registration, of eligible issuers, and how issuing should be conducted. There are distinct differences between the domain protocols, as for example some issue certificates weekly, as in the case of Norway, and some bi-monthly or monthly, as in the case of Germany and Spain. [46, 47, 48, 49]

The Association of Issuing Bodies (AIB) is appointed by the EU, and it is responsible for overseeing, developing, and implementing the rules and guidelines that each domain must follow for their certificates to be approved as a GOs. The AIB manages the EECS, which encompasses the entire process of GO production, trade, and consumption. The AIB audits the trading system, keeps statistics and manages the AIB trading hub. The trading hub serves as a centralized platform for cross-domain trading of GOs, where all certificates moving between domains must pass through before it can be registered as the property of the receiving domain.

A short overview of cross-domain trading can be as following: [4]

1. The GO holder initiates a trade by contacting the issuing body in its respective domain.
2. The issuing body contacts the AIB hub, which registers the trade.

3. The GO is now in transit, and can either be canceled or approved.

4. If approved, the GO is forwarded to the receiving domains issuing body.

5. Finally, the receiving issuing body forwards the GO to the final receiver.

Until recently, the trading of GOs primarily occurred through peer-to-peer transactions or a limited number of national auctions. Countries such as France, Portugal, the UK, Italy, Croatia, Hungary, and Slovakia had established their own auction platforms. [50] However, trade beyond these regions was facilitated by third-party services like Commerg [51] and the Green Power Hub [52], or through direct transactions among major market participants.

However, in September 2022, the first pan-European spot auction took place, marking the establishment of an organized and unified marketplace [53]. This auction was open to countries participating in the EECS, and it successfully facilitated the trade of GOs equivalent to 19 171 MWh of energy. Subsequently, an October auction followed, where 69,326 MWh were traded [54]. The price range for the GOs in these auctions was between 5.76 and 5.89 e/MWh, reflecting market dynamics and demand.

The exchange of GOs within and outside of the EECS between 2019 and 2022 is presented in Figure 7 as two node networks. The networks depict the net flow of GOs, meaning that although many countries engage in reciprocal trading this will cancel out, and only the difference in imports and exports are presented. The nodes in the network are labeled with standard country abbreviations and encompass both European and non-European countries. In Figure 7a, the node sizes are logarithmically scaled based on each country's GO import, while in Figure 7b, the node sizes are logarithmically scaled based on each country's GO export.

Interestingly, the largest importer of GOs during this time periode was Great Btritain. This is despite not being a member of the EU or the European Economic Area (EEA). However, they do recognize GOs as an alternative to their own REGOs [55]. Contrary, Norway stands out as the largest exporter of GOs, exporting approximately 21 times more than the second-largest exporter, the Netherlands. While Norway primarily supplies European countries, it also exports GOs to several non-European countries including as Peru, Australia, Brazil, and Chile.

Trading GOs across domains is a process consisting of several third party actors and convoluted and bureaucratic transits. A cross-domain GO trade includes a minimum of five bodies, starting from the GO holder, then moving to the domain's issuing body, then through the AIB hub to the receiving domain's issuing body, before it finally is received by the customer. Throughout this transit, the GO is not interactable, and is de facto being held in escrow by third parties. By default a canceled trade returns the GO to the original holder, and an approved trade extends the GO to the customer. Long and winded transit means that there is a significant time frame where the accounting of GOs and the digital holding do not correspond. Rather, GO accounting must be done on the assumption that trades are valid and proper.

In Figure 8, a flowchart of the current trading architecture is presented. This flowchart shows the entities involved in a single cross-domain trade, but the full architecture would include similar structures for all domains. The flowchart shows the flow of GOs from it is issued to it is traded. The blue components of the flowchart are GO flows, with the newly opened Epexspot GO auction in the middle. Here, standard GO prices are set and trades are planned, however

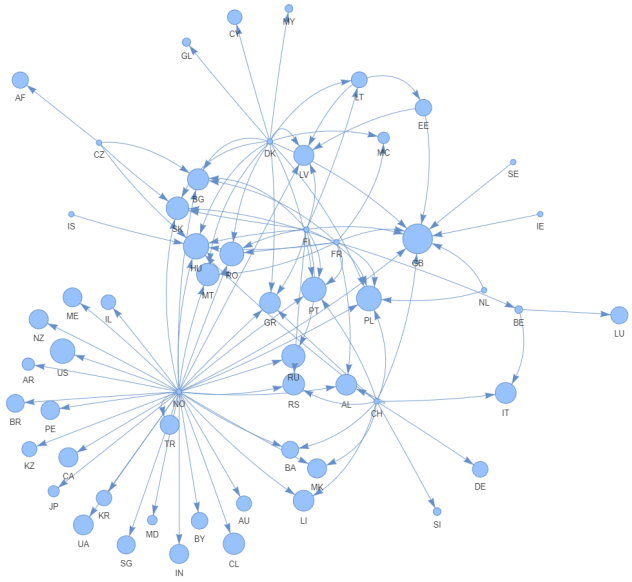

(a) Size of country node represents GO consumption.

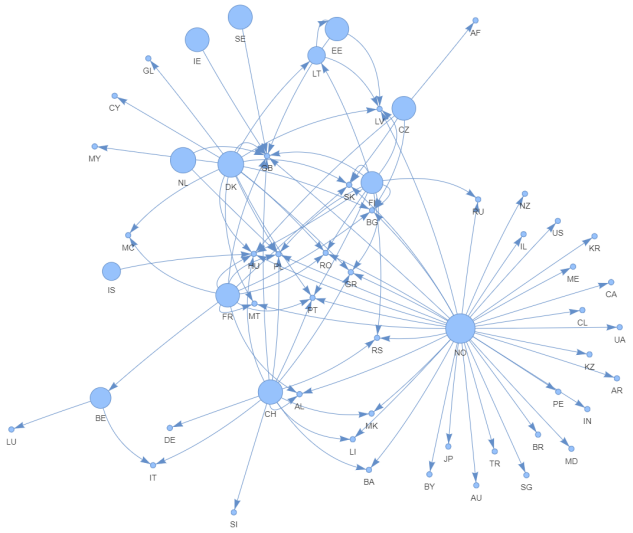

(b) Size of country node represents GO provision.

Figure 7: Trades of GOs from 2019 to 2022. Data from [56].

the concrete transfer of GOs goes through the longer process which includes the issuing bodies of each domain, and the AIB hub. The concrete GO transfer process is represented by thick black arrows. The orange blocks are the bodies that may hold the GOs at some point during the transfer process.

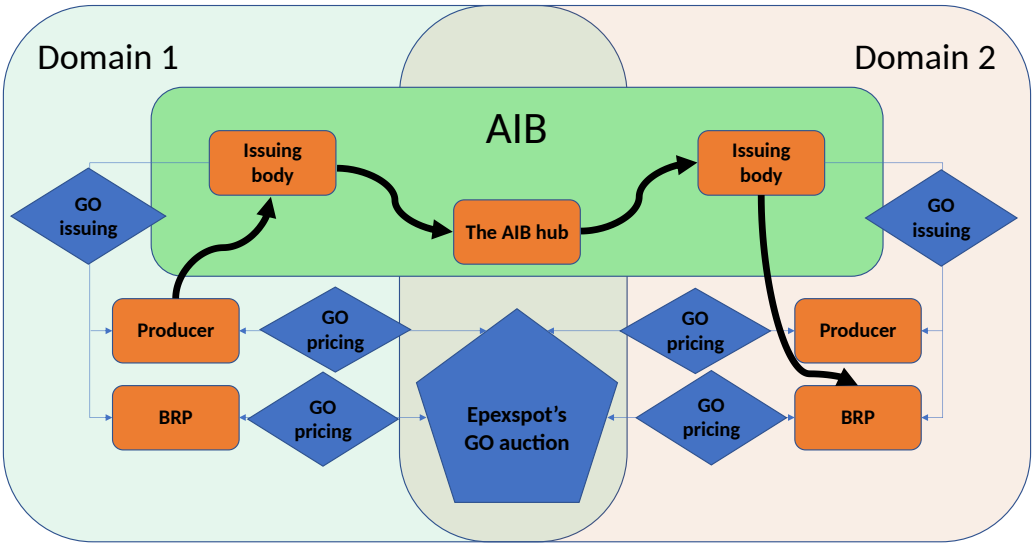

Figure 8: A flowchart of the current GO trade architecture.

3 Distributed Ledger Technologies and Blockchain

3.1 Distributed Ledger Technologies

DLTs encompass various technologies, primarily blockchain, but includes among others also directed acyclic graphs, hashchains and holochains. [57] These technologies share some common traits, such as the ability to store data in a decentralized manner and facilitate peer-to-peer communication. They employ cryptographic techniques to ensure immutable storage and consensus mechanisms to establish the legitimacy of information added to the ledger [57]. DLTs can be classified into four subcategories based on two criteria: permissioned/permissionless and private/public, where each combination of these criteria have unique characteristics:

1. Permissionless and public [4]

   Permissionless and public DLTs are the most widespread combination of criteria. Publicly traded cryptocurrencies fall under this combination, including for example Bitcoin and Ethereum. Permissionless and public DLTs are open for any actor wishing to take part, and all nodes are allowed to both partake in transactions, and act as validators of other transactions and blocks. This combination is unique in that it is completely trustless and decentralized, and that all nodes have equal privileges. [58] [57]

2. Permissionless and private [4]

   Permissionless and private DLTs are not uniformly agreed upon as feasible. In [58], the authors claim that the combination of permissionless and private itself is an oxymoron, as private DLTs require a verifiable identity for participation, while permissionless implies no access-control. However, other authors define the roles slightly different, and in [57] a permissionless/private DLT is one where entry to the network is limited, but roles within the network are equal. This means that only allowed actors can participate in the network, but once admitted, everyone can both suggest data to be added to the ledger, and help validate the legitimacy of the data. [58] [57]

3. Permissioned and public [4]

   Permissioned and public networks, also sometimes called hybrid networks, are networks where there are no restrictions on who can see and access the ledger, but other actions may be restricted. On a hybrid DLT, participants could need permissions to transact, validate or suggest blocks for validation. After Ethereum's change from Proof of Work (PoW) consensus to Proof of Stake (PoS) consensus, they have become a permissioned and public DLT, as everyone can suggest transactions, but only participants that are sufficiently invested can validate transactions and consolidate blocks. [58] [57]

4. Permissioned and private [4]

   Permissioned and private networks are the most restrictive, as they require both access and acting permissions. Such a network lends itself well to organizations and enterprises where definite roles are set, and only organizational personnel should have access. For example, an enterprise may want to have specific rules for which transactions should be prioritized when adding transactions to the ledger, and therefore may have only a handful

of nodes proposing transactions for validation. [58] [57]

3.2 Blockchain

Blockchain, the most well-known form of DLT, is often closely associated with digital currencies like Bitcoin. However, its potential encompasses a broader array of use cases. In addition to being employed in the finance sector, it has been explored in areas such as the energy sector, the healthcare sector and alongside intellectual property.[59, 60, 61] In general it serves as a decentralized and synchronized ledger, enabling direct asset transfers between clients without the need for collateral. Any information that benefits from being stored transparently and securely could reasonably be stored on the blockchain. At its core, a blockchain simply stores information and informtion changes, but ensures that this information is distributed and synchronized. This makes blockchains intrinsically resistant to faults by not operating with a single trusted point.

The figure in Figure 9 provides a simplified timeline of the process that occurs when a transaction occurs on the blockchain. The term is widely used for any change to the information stored on the blockchain, and while it comes from the transaction of Bitcoins, it more generally means the change of a blockchain state. This is a result of blockchains becoming more generalized over time, shifting from being a purely financial instrument, to becoming a general information ledger.

The process depicted in the figure begins with someone initiating a transaction. Initially, this transaction is stored temporarily as a non-confirmed floating transaction. This temporary transaction storage is refered to as the transaction pool. Following, a miner, staker or other validator includes the transaction in a proposed block. Once the validator is satisfied with the assembled block, they broadcast it to the rest of the network. The block is then subjected to verification by other validating nodes within the network. If the block is deemed legitimate, it becomes the latest extension of the blockchain, and if the block is successfully added to the blockchain, the state change associated with the transaction is considered approved, and the transfer is completed.

The process of gathering, broadcasting, and validating a new block of transactions may vary across different blockchains but is encapsulated within the blockchain's consensus protocol/mechanism. Consensus protocolls are the mechanism that ensures the legitimacy of blocks and transactions, and usually requires a majority consensus to alter the blockchain. More details about consensus mechanisms can be found in subsection 3.4.

The design of a blockchain can vary depending on its specific use-case. It can be designed to minimize the need for user trust, enabling clients to benefit from the blockchain's traits without depending on the trustworthiness of others. Alternatively, it can also incorporate a certain level of inter-participant trust, which can enhance the speed flexibility of operating on the blockchain, and minimize the size of the blockchain itself. The abilty to create a completely trustless system is made possible by the decentralized nature of blockchain networks, where multiple nodes collaborate and rely on consensus mechanisms to authenticate, verify, and validate asset transactions. The various consensus mechanisms offers blockchains that range from fully trustless networks to networks with a central trusted authority.

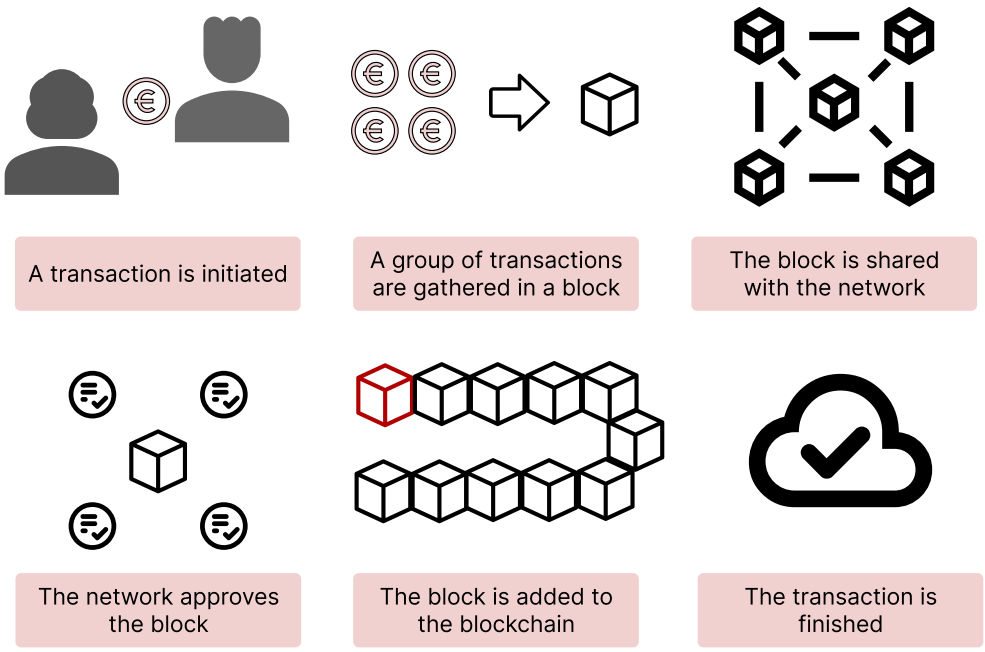

Figure 9: An example of the general order of procedures associated with conducting a transaction on a blockchain.

Network with nodes ranging from trusted to trustless generally comes with a tradeoff in transaction speed, storage space used and resource consumption. More generally, this tradeoff is considered a tradeoff between scalability, security and decentralization, and was first introduced by the founder of Ethereum, Vitalin Buterin. [62] In the trilemma, decentralization refers to the ability for any and all parties of the network to have an equal share of information, and to equally trusted by all other parties. The scalability of a blockchain refers to its ability to operate as more nodes are connected, and information is shared widely. Scalability has been a problem with completely trustless and decentralized blockchains, as the consensus mechanism keeping them secure and consistent has demanded increasing amounts of resources. For example, as the Bitcoin network has grown, the computational power needed to mine an average block has grown as well, and today the energy use from this could be considered unsustainable. [63] A naive solution that would allow blockchains to be both scalable and decentralized is to decrease the security of the network, which is the third pillar of Butterins trilemma. This is rarely done, but could be considered on private networks where all paries are known and trusted. An example of this is the use of the RAFT consensus mechanism over the proof of Byzantine Fault Tolerance (pBFT) consensus mechanism. These consensus mechanisms are both utilized in private networks, but the RAFT mechanism is only crash resistant, while the pBFT mechanism is also tolerant to some malicious actors. [64, 65] See subsection 3.4 for more.

Generally, newer blockchains have made progress in improving scalability, without severely sacrificing the decentral nature and security of the blockchain. The IBM backed blockchain Hyperledger Fabric has achieved higher transaction speeds and reduced energy consumption, but its consensus mechanism relies on either Proof of Identity (PoI) or Proof of Authority (PoA) to ensure security. [66] The study in [67] suggests that any consensus mechanisms based on PoI or PoA may compromise decentralization to establish distinct and reliable identities.

Implementations of unique digital identities can be undermined by external actors using multiple identities, unless a governing body or trusted authority verifies the one-to-one linkage between users and digital identities. In other words, blockchains using PoI or PoA consensus mechanisms intrinsically cannot be both fully decentralized and secure.

Another instance highlighting the Scalability Trilemma can be seen in the proposed "solution" presented in [68]. They suggest a committee-based PoS approach similar to the one planned to be introduced through the sharding update launched by Ethereum (see subsubsection 3.6.1) [69]. Although this solution is generally regarded as secure and scalable, it does face reduced security compared to standard unclustered, or "un-comittied" chains. As detailed in [69], the threshold for chain ownership required for an attack is 51% in unclustered networks, but it decreases to approximately 30-40% in randomized committee-based chains. These chains do however, offer significant advantages in terms of large-scale transaction speeds when compared to standard unclustered chains.

3.3 Tokens in Blockchain

Tokens play a profound and multifaceted role in the blockchain ecosystem. While they are commonly recognized as the tradable unit behind cryptocurrencies, their role in generalized blockchains is more varied. Tokens have the pivotal role of representing and enforce scarcity within the chain, enabling divisibility, transferability, and tradability [70]. Tokens facilitate transactions among participants, and their values can be recorded as gains or losses on the digital ledger. Since tokens are inherent to blockchains themselves, they can be used create incentive structures where block validators can be rewarded without impacting other participants negatively. Instead, the rewards are obtained through a regulated token dilution rate that is shared among the network [71] In [70], a collection of eight use cases for tokens in blockchain has been compiled from various sources. These use cases not only emphasize the necessity of tokens in public cryptocurrencies but also highlight unconventional applications that broaden the spectrum of blockchain technology. The compiled list of these use cases includes [4]:

- Currency

  Using the token as a representation of value, allowing it to stand in as a monetary digital object.

- Validation incentive

  Representing a reward for validation of blocks, and therefore acting as an incentive for users to invest resources into building the network.

- Usage incentive

  By being located on the blockchain, users are incentivized to participate in order for their tokens to gain value.

- Network participation acceleration

  Early adopters can be rewarded with higher token incentives.

- Governance and management

  Allowing for rules and protocols to be enforced on the blockchain by the means of token costs and rewards.

- Representation of asset ownership

  By representing the ownership, or right to some non-monetary asset. For example a unit of energy. [72] [73]

- Profit-sharing

By delegating dividends or profit on the basis of token-holdings.

- Funding instrument

As an instrument for public funding though either Internal Coin Offerings (ICOs), or crowdfunding. [14]

[70]

Fungible tokens are the most common type of tokens in the blockchain space. A token is considered fungible when it is interchangeable with other tokens of the same type, and all tokens of that type are considered to have the same value. [74] For instance, in the case of the standard Ethereum token, the ETH, it is at all times interchangeable with any other ETH. If one ETH is exchanged for another, the parties involved will have ann equivalent token holding.

The alternative to fungible tokens, the Non-Fungible Tokens (NFT), has gained significant public attention, especially in the realm of digital art and assets. The attention followed after a surge in NFT sales on the fiat currency market between late 2020 and early 2022 [75]. However NFTs are not exclusively associated with digital art, instead they originated from the ERC-721 Ethereum token standard as unique alternatives to fungible tokens. [76] Unlike fungible tokens, where the exchange of one token for another doesn't result in any change, NFTs introduced by the ERC-721 standard are unique and traceable. They allow for the ownership of a distinct asset to be recorded on the blockchain, making it verifiable by anyone with access to the blockchain. Each NFT has a specific token ID that can be linked to a designated wallet address, granting the user exclusive ownership of that ID. Some suggested use cases for NFTs include collectibles, access keys or tickets. [77]

A third alternative has also emerged in the blockchain space. The Semi-Fungible Token (SFT) is somewhere between the NFT and the fungible token. It allows for grouped tokens, meaning that groups of tokens with similar characteristics can be created and IDed. Any tokens within an ID group will be exchangable and fungible with one another, but each group will be uniquely represented. The Ethereum ERC1155 standard implements the ability to create SFTs, such several token groups can be created and traded with the same smart contract. [78]

3.4 Consensus Mechanisms

The key characteristics of blockchains rely heavily on robust consensus mechanisms. Key attributes like resistance to double spending, immutability, and verifiability are all contingent upon consensus mechanisms that effectively maintain synchronization and agreement across the chain. Consequently, numerous consensus mechanisms have been devised and investigated in various blockchains, including private/public and permissioned/permissionless architectures. [79] While there are multiple approaches to categorizing these consensus mechanisms, a commonly used classification is proof-based and voting-based.[80]

Proof-based consensus mechanisms, including popular ones like PoW and PoS, are commonly referred to as Proof-of-X (PoX) mechanisms. These mechanisms require participating nodes to provide evidence or proof of their commitment to the blockchain. Once this proof is provided, nodes can publish a block to the blockchain, which is then validated and adopted by the rest of the network. The purpose of requiring proof is to discourage malicious actors from attempting

to manipulate the blockchain. [81] In the case of PoW, the proof of dedication involves solving a computationally intense problem. This process requires a significant investment of energy and computational resources. The provided proof is a "nonce," a number that, when combined with block-related data, produces a hash value lower than a predetermined difficulty target. [79] Other PoX mechanisms, such as PoA, rely on a set of predetermined validators who possess confirmed authority. In PoA, the proof lies in the trustworthiness of these validators. While PoA can be used in public networks, it is primarily deployed in permissioned and private blockchain networks. [79]

Voting-based consensus mechanisms are typically implemented in permissioned and private blockchain networks, where the nodes are reliable and known entities. This is because voting-based mechanisms rely on a predetermined number of nodes to achieve a consensus threshold. [80] If the number of nodes in the network fluctuates without the knowledge of other nodes, it becomes challenging to establish a consistent threshold for reaching consensus. However, under stable conditions, a voting-based consensus mechanism can be efficient and resilient to faults. [80] Two examples of such mechanisms are pBFT, as described in subsubsection 3.4.2, and the RAFT consensus algorithm. [64] Both mechanisms are designed to achieve consensus while being resilient to some faults. RAFT is resistant to crashed nodes, meaning it can continue functioning even if some nodes become unresponsive or fail. pBFT is comparatively resistant to both crashed nodes and malicious voter-nodes, therefore being more eligible for environments with less trust.

When a consensus mechanism determines a legitimate block, there is a possibility that another legitimate block may be created by a different node before the blockchain is fully synchronized. This situation leads to a fork in the blockchain, where two valid chains exist simultaneously for nodes to work on. [82] The typical approach to resolving a fork is to wait for one of the chains to become longer than the other. The longer chain is then considered the main chain, while the redundant block is referred to as an "ommer" block on the Ethereum network or an "orphan" block on the Bitcoin network. On the Bitcoin blockchain, nodes that have mined but are not chosen are simply discarded, and the miner receives no reward. Ethereum has implemented a fairer system where the creator of the ommer block is rewarded with a fraction of the staker's reward. It's important to note that forks resulting from ommer and orphan blocks should not be confused with soft forks or hard forks. Soft forks occur when updates are made to the blockchain that do not require the entire userbase to update for the new blocks to be considered valid. Soft forks can be accepted by the blockchain once a majority of the nodes have adopted the update. On the other hand, hard forks are the result of disruptive updates that fundamentally change the way transactions are handled. These types of forks require all nodes wishing to operate on the main chani to implement the update. [83]

3.4.1 Proof of Stake

PoS has emerged as a popular alternative to PoW, offering a different approach to achieving consensus. In PoS, participants known as "staker" peers voluntarily stake a certain amount of tokens with the aim of validating new blocks. The terminology used in PoS draws inspiration from the world of betting, as block confirmation resembles a majority bet. With each block, a leader is selected from a pool of eligible validators, who then proposes a block for inclusion in the blockchain. To qualify as a potential leader, a peer must a minimum of tokens, serving as

evidence of their commitment to maintaining the integrity of the blockchain. The underlying idea is that individuals who have a significant stake in the blockchain are less likely to engage in actions that could compromise its integrity.

PoS leader selection has been explored through various methods, where two major methods are "stake size" and "stake age." Stake size-based leader selection straightforwardly chooses the leader with the largest stake, but this approach has inherent issues of wealth concentration and high barriers for newcomers. A more refined method is the "stake age" approach, which takes into account both the stake size and the stake holding time. [84] When a leader is chosen, their staking age is reset to zero, allowing other candidates from the validator pool to have a chance. This enables even small stakeholders to potentially become leaders if they stake their resources for a sufficient duration. Although the stake age method partially addresses the "rich get richer" problem, it is not entirely eliminated. See subsubsection 3.5.2 for a better explenation of the "rich get richer" problem in PoS.

Once a leader has been chosen and proposes a new block, they will stake their own tokens on the block's legitimacy. Following, the remaining validators in the pool will participate in a voting process to either confirm or deny the proposed block, with their voting power determined by the size of their stake. If a majority of the validators validate the block, the proposing leader will receive a token reward proportional to their stake. However, in the basic version of PoS, there is a concern known as the "nothing at stake" problem, where there are no consequences for staking on invalid blocks. To address this issue, many PoS blockchains have implemented mechanisms to punish such behavior. In Ethereum, for example, if a block is deemed invalid, the proposing leader will lose their stake and be removed from the validator pool. [85]

Unlike for PoW, PoS does not rely on a computationally heavy and time consuming validation mechanism, which may allow it to reach significantly faster transaction speeds. In fact, PoS and hybrid PoS-PoX blockchains have achieved transaction speeds that surpass PoW mechanisms by several orders of magnitude. For example, the Algorand blockchain has demonstrated an average transaction speed of 800-1000 transactions per second, and optimized Ethereum claims to potentially be able to hit 5200 transactions per second. [86, 87] Still, this is slower than other consensus mechanisms, such as the Proof-of-History used by Solana, which consistently reaches 4000+ tps, and claimed to be able to reach up to 65 000 transactions per second during testing [88, 89].

3.4.2 Practical Byzantine Fault Tolerance

The pBFT protocol, represents a modified version of the general Byzantine Fault Tolerance algorithm. It offers the advantage of asynchronous operation and higher speeds [90]. Byzantine Fault Tolerance, originating from the Byzantine empire, predates digital technologies and has has been applied in a variety of circumstances. [91] Initially developed in 1999, the pBFT algorithm served as a tool for secure and verifiable communication during the early stages of the internet. Over time, research and developmental efforts have been dedicated to pBFT, particularly in the context of a consensus mechanism for DLTs. [65] Although pBFT has been a subject of significant research and is often mentioned as a potential candidate for an effective blockchain consensus mechanisms, its practical implementation has been somewhat limited [65]. However, there are concrete examples of blockchains implementing pBFT as the

consensus mechanism, most notably Hyperledger Fabric. [92]

The pBFT consensus mechanism follows a three-step process, where each phase is validated by two-thirds of the network nodes, including the node itself. In a network with a chosen number of nodes, denoted as "n", consensus can be reached as long as the number of faults, denoted as "f", is less than or equal to $f = \lfloor \frac{n-1}{3} \rfloor$. The algorithm consists of the following three phases [4]:

1. Pre-Prepare phase

   In the pre-prepare phase, the primary node broadcasts a message to all backups in the network for them to validate. The backups note the message in its log.

2. Prepare phase

   Given that the pre-prepare message is validated by the backup node, the backup nodes broadcasts a prepare message to all other replica nodes in the connected network. If the prepare message is the same as the pre-prepare message from the primary node, the prepare message is logged by the receiving nodes. A node is considered prepared when it has $2f$ the number of validated prepare messages in its log.

3. Commit phase

   After a node is properly prepared, it sends a commit message, confirming that the block is valid. When the primary receives $2f$ commit responses, so that it has $2f + 1$ validations, including its own, the block is committed.

The complete implementation of the pBFT algorithm can be found in [90]. While pBFT provides security against a limited number of faulty nodes, it relies on trustworthy and secure primary nodes. Its vulnerability to malicious consensus leaders is a well-known drawback, which is why pBFT is predominantly utilized in permissioned and private blockchain networks like Hyperledger Fabric. Some solutions to malicious leaders have been developed, but then there exists a trade-off between transaction speed and resilience against malicious actors [93]. However, when the consensus leader can be trusted, which is likely on a permissioned and private network, pBFT has demonstrated impressive transaction speeds, reaching up to 21,000 transactions per second in the case of Hyperledger Fabric [66].

3.5 Common Problems and Security Concerns

3.5.1 The 51% Attack

Consensus mechanisms are designed to make attacks unfeasible, however over time unexpected attack angles have been discovered. The most commonly mentioned security concern is the 51% attack. It is mainly talked about in the context of PoW or PoS consensus mechanisms, as they have consensus mechanisms with consensus power that scales based on either the computational power or stake size. An attacker would need to hold more than 50% of the hashing power, or 50% of the total blockchain stake over an extended period of time, which is why the attack is hard to execute, especially on larger blockchains. For a PoW blockchain, a 51% attack could be used to double spend tokens by: (1) Spending tokens normally on the blockchain. (2) Delegating the >50% hashing power to mining and adding blocks to a shorter fork where the

spent transaction is not included. (3) Since the attacker controls more than 50% of the hashing power of the network, the alternative fork will eventually become longer than the legitimate chain, and the rest of the network will adopt this now longer chain as the legitimate chain. (4) The previous transaction made by the attacker is now not recorded on the ledger, and the tokens can be spend again.

When Bitcoin was first launched in 2009, it was thought impossible to achieve a 51% attack, but the development of hashing equipment has facilitated some attacks on small and medium size blockchains. Some examples are [4]:

- In January 2020, a hard fork of Bitcoin, Bitcoin gold, experienced a 51% attack. It is estimated that approximately 72 000 $ in bitcoin gold tokens were double spend during the attack, and that the attackers had a hashrate of 2.53 MH/s. [94]
- In July 2019, an Ethereum spin-of chain called Expanse was attacked. The attack removed 63 blocks and replaced them with 64 new blocks. Only around 12 $ were double spent. The Expanse network used the Ethash algorithm, which was the PoW mining algorith used by Ethereum before the switch to PoS. [95]
- In July 2019, Litecoin cash was attacked. It uses a hybrid PoW/PoS consensus algorithm that interlaces PoS blocks with PoW blocks, making it more difficult to produce reorganized longest chains. Around 4600$ were double spent during the attack, but the attackers were unable to produce longer alternative chains due to the hybrid PoW/PoS. [96]

A blockchain's resilience to 51% attacks relies largely on the size of decentralized and distributed nodes. However, an increase in mining and staking pools have raised concerns regarding attack resistance, especially when looking at Bitcoin and Ethereum. At the time of writing, the two largest Bitcoin mining pools have 58% of the hashing power of the entire network, and the three largest have 68.8%. [97] While these pools are not single nodes with large amounts of computational power, they have been delegated control in order to mine efficiently and hash different nonces and transnational data. These pools have the coordination to create a 51% attack should they wish, as has been seen in 2019, when a vulnerability in Bitcoin cash (BTH, not BTC) led to it being exploited for 1.35 million $, but the pools executed a benevolent 51% attack to restore the chain to its non-fraudulent state. [98]

3.5.2 Long Ranged attacks and Stake Bleeding

The PoS consensus mechanism has brought attention to a compound problem known as the "nothing at stake" issue, which gives rise to concerns such as long-range attacks and stake bleeding. The nothing at stake problem arises from the absence of a cost associated with reaching a consensus decision in simple PoS systems. Stakers are required to stake their tokens on a specific chain to claim the rewards associated with confirming a new block. However, if there are no consequences for staking on the incorrect chain, rational stakers would be incentivized to stake on every available chain [99].

Stake bleeding occurs when a validator voluntarily gives up the ability to propose blocks to the chain. When coordinated, this action can impede the progress of the chain, providing an

opportunity for alternative chains to catch up. By considering a scenario where a 10% stake ownership grants the ability to propose 10% of the blocks added to the chain, a stake bleeder could reduce the speed of chain growth to 90% of its normal rate by consistently forfeiting the chance to create new blocks [99].

The combination of the "nothing at stake" property and stake bleeding on the main chain introduces the possibility for a malicious actor to construct an alternative side chain that could gradually surpass the length of the legitimate main chain. To achieve this, the attacker would follow a two-step approach. First, the attacker would deliberately forfeit every opportunity to propose blocks on the main chain, intentionally slowing down its progress. As time passes, the attacker's proportional stake on the legitimate chain diminishes, reducing the impact of the forfeited blocks. Secondly, the attacker would work on a hidden side chain where they are the sole staker, allowing them to collect all the rewards associated with staking and validating blocks. Since the secret side chain does not include any forfeited blocks, it gradually catches up to the legitimate chain, eventually surpassing it in length. Once the side chain reaches a sufficient length, the attacker presents it to the network as the main chain, which, being longer, is accepted. As a result, the attacker receives all the staking rewards tied to the newly accepted chain, while the rewards tied to the legitimate chain are disregarded [99].

To mitigate the risks of long-range attacks, one approach is to implement checkpoints within the blockchain. This involves introducing "universal" blocks that are hardcoded and required within the legitimate chain, as implemented by Ethereum [100]. These nodes serve as backward boundaries, limiting the starting point of a potential side chain and reducing the timeframe available for a long-range attacker to construct an alternative chain. Another method involves storing the state of the chain in a separate blockchain. As proposed by [101], one possibility is to store chain states on the bitcoin ledger, eliminating the necessity for a central authority to determine the inclusion of blocks as "universal." This decentralizes the process and enhances the security of the consensus mechanism.

3.6 Ethereum

Ethereum is commonly recognized as the second largest blockchain, although it is challenging to establish precise metrics for measuring blockchain "size." As of now, Ethereum holds the second position in terms of market cap [102]. Ethereum made an early entry into the cryptocurrency realm, publishing its whitepaper in 2013 and mining its genesis block on July 30th, 2015 [103]. While other blockchains had built upon Bitcoin's architecture, Ethereum aimed to differentiate itself by separating the blockchain from Bitcoin's transactional nature. Instead, Ethereum allowed users to deploy the entire backend of applications on the blockchain, enabling the storage of general variable states [104]. Although often referred to as storing transactions, Ethereum transactions contain broader state changes than those associated with traditional transactions.

The Ethereum blockchain is inherently different than purely fiscal blockchains, in that it can store any type of information that can be represented as a state of bits. These states are changed and managed by the Ethereum Virtual Machine (EVM), which is an emulated machine that uses the blockchain ledger as its permanent storage, and a local machines storage as volatile storage. [105] Within a normal computer, the ledger would be analogous to the hard drive, and

the local storage would be analogous to the RAM. The EVM is a stack-based machine, and operates the stack based on given opcodes.

Alongside the EVM, the Ethereum foundation has developed and released a fully Turing complete language called Solidity. Compared to earlier attempts at blockchain code, Solidity implements conditional statements, loops and multi-stage transactions through states [106]. Later blockchains have been developed with similar developmental freedom, but Solidity was unique at release. The Solidity language is inspired by C++, but has additional functionalities specifically targeting blockchain related mechanisms such as paying to and withdrawing from contracts, and altering the blockchain state. It also has the inherent ability to access the data of interacting transactions, such as the value or sender of the transaction, or a message sent with it. [107] Before Solidity code is run, it must first be compiled down to opcodes and bytecodes. Opcodes are low level operations compiled down from higher level solidity code, and they are both human and machine readable. The opcodes are then represented with a hex value for the machine, which are called bytecodes. [108] By compiling the language rather than interpreting it, solidity developers gain the ability to discover and improve type-errors and other compile-time errors before the contract is launched onto the blockchain. This prevents wasted gas fees otherwise incurred by deployed erroneous contracts.

Gas fees refer to the relatively small transaction fees incurred when deploying or interacting with smart contracts on the Ethereum network. These fees are determined by several factors, including the operations performed by the EVM, the extent of the state change made on the blockchain, and the nature of the transaction itself. For example, gas fees can vary depending on whether the transaction involves contract deployment, transferring funds between addresses, modifying specific states, or simply accessing information on the blockchain.[109, 110] Given the impact of gas fees on transaction costs, developers should optimize their smart contracts to minimize gas consumption, especially when dealing with contracts that are expected to be frequently called. Ethereum stores data in 256-bit chunks, even if the actual data size is smaller. However, through a technique called "stacking," multiple variables can be efficiently packed into a single 256-bit package by arranging them in the correct order. [111] This can be gained by developers by defining variables of different sizes in orders that add up 256-bits.

3.6.1 Ethereum and the Switch to Proof of Stake

Since its launch in 2015, the Ethereum blockchain has undergone significant updates. Initially, Ethereum operated with a PoW consensus mechanism, however, due to the high energy consumption and environmental impact associated with PoW, the Ethereum Foundation recognized the need for a more efficient and sustainable approach. [112, 113] As early as 2014, the transition to a PoS consensus mechanism was proposed. Originally the update was referred to as Serenity and later as Ethereum 2, however these names have since been deprecated as the PoS deployment has been divided into several major updates.

’ The first milestone in the transition to a PoS consensus mechanism was the introduction of the Beacon Chain. On December 1st, 2020, the Beacon Chain was launched as a parallel chain operating alongside the mainnet Ethereum chain [103]. It served the purpose of testing and ensuring the proper functioning of the new consensus mechanisms. The Beacon Chain introduced the concept of a "Consensus layer," which distinguished it from the execution layer responsible

for transaction execution and smart contracts. While the Beacon Chain implemented the new consensus layer, it did not modify the execution of transactions and contracts. [114]

The second milestone in the transformation of the Ethereum blockchain was the Merge, which was implemented on September 15th, 2022. Its objective was to merge the Beacon Chain with the Ethereum mainnet, effectively transitioning the consensus mechanism from PoW to PoS. According to official Ethereum documentation, the Merge successfully reduced the chain's power consumption from over 5 GW to less than 3 MW, meaning a decrease of around 99.95%.[115] Following the Merge, the inter-block time is now fixed, with blocks being approved on a predefined schedule every 12 seconds (or a multiple of 12 seconds if blocks are skipped). The introduction of PoS replaced miners with stakers, requiring validators to stake 32 ETH to participate. While the significant entry cost is intended to ensure the commitment and security of the network, it does pose challenges to the decentralized nature of the chain, often favoring staking pools over individual stakers. It is worth noting that the concentration of mining power in large mining pools is not a problem unique to PoS, as it has also been observed in PoW systems.

Originally a third major milestone was planned to improve the transaction speed of Ethereum. It was the inclusion of sharding, a way to split the blockchain in to smaller chains such that work could be parallelized. However, due to efficient 2. layer options, the sharding update has been discontinued [3]. For example, the addition of transaction rollups has allowed users merge transactions outside the blockchain, and then execute them as a single transaction on the blockchain. This limits the fixed costs related to executing transactions, and therefore potentially the gas, space and time costs. [116] It is estimated that optimal use of rollups will allow for rougly 78 000 transactions to be bundled in a single block, giving transaction speeds of roughly 5200 transactions per second. [87] Therefore, an update aiming at optimizing interoperability between 2.layer solutions and the Ethereum blockchain is under construction. This update is named Danksharding, and is estimated to release part by part. The first update is expected to happen sometime during 2023. [117]

4 Methodology

For this thesis, a two-part job has been done. The main work consists of creating, testing and evaluating a marketplace for GOs based on blockchain technologies, but to accommodate the certificates employed on the marketplace, there is also a co-simulation used to create realistic time-series data for power production. The on-chain marketplace and the power co-simulation forms the basis for a physical layer simulation and a cyber layer simulation. This section will explain how the simulation and co-simulation are used to create a node network where each node will interact with each other over the marketplace. An overview of the entire methodological structure can be seen in Figure 10. Note that there are two flows of data to the cyber layer, one flow for the energy produced, which in turn determines the amount of certificates created, and one flow for the energy consumed, which will determine the certificates consumed by each node.

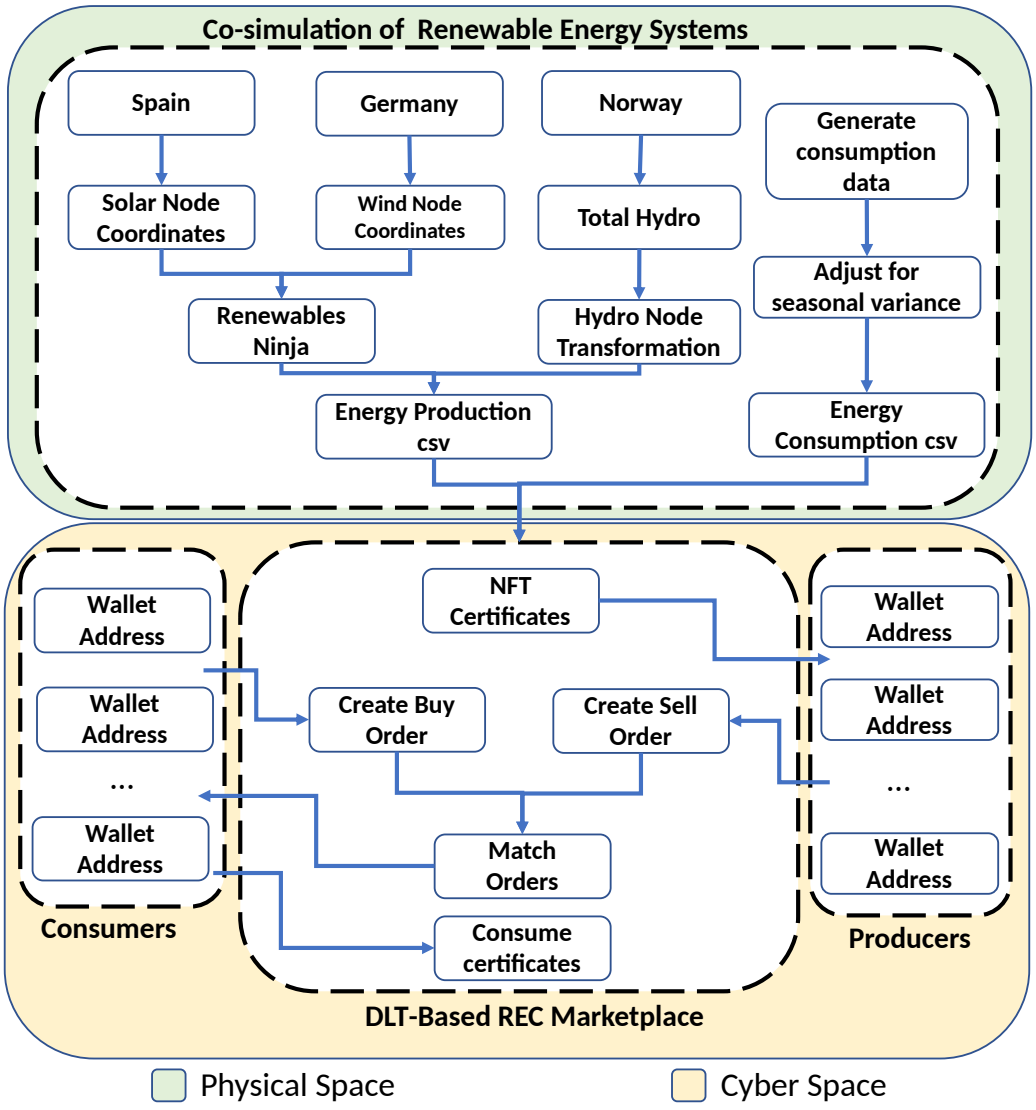

Figure 10: An overview of the simulation parts, with a distinction between cyber space components and physical space components.

4.1 The Trade Node Network

As a co-simulation to the certificate marketplace, an energy production simulation is included. As GOs are predominantly a European certificate, the co-simulation focused on European renewable energy production. The aim of the co-simulation is predominantly to provide some time-series data that can found the base of the certificate production implemented on the marketplace. The co-simulation simulates clean energy produced in Germany, Spain and Norway, specifically wind, solar and hydro energy. In the co-simulation, each country consists of five nodes producing and consuming certificates, but a real world implementations would have many more producers, consumers and prosumers for each country, and a wider variety in the types of certificates added to the blockchain.

4.1.1 Solar Power Simulations

The spainsh nodes are simulating solar energy production through the use of an online tool named Renewable Ninja. Renewable Ninja is well documented and aknowledged tool for wind and solar simulations, which creates accurate solar power predictions for a given time frame and location. [118] This means that the solar predictions are made using a mix of irradiance data and temperature data. The temperature data is taken from the Modeling-Era Retrospective Analysis for Research and Applications (MERRA)-2 dataset provided by NASA. This dataset is notable in that it provides global athmosperic data, including for example temperature measurements. Renewable Ninja uses irradiance data either directly from the freely available satellite derived dataset CM-SAF SARAH, or derives it from MERRA-2's atmospheric data. The former dataset is only available in Europe, while the latter is available globally. [119]

To compute the PV-power, both an ambient temperature and the irradiance is considered. The MERRA-2 dataset provides global surface temperatures used for as the ambient temperatures, and the panel irradiance can be calculated from site specifications and the SARAH irradiance data. The panel's irradiance consists of two parts, the direct irradiance from the sun, and any diffused irradiance form nearby reflective surfaces or cloud diffusion. The two equations Equation 1 and Equation 2 shows how the direct and diffused irradiances are calculated. The direct irradiance is subject to Equation 3, which accounts for the the plane incident angle between the direct irradiance from the sun and the panel.[119]

$$I_{dir,p} = \frac{I_{dir,h} \cdot cos(\alpha)}{cos(\frac{\pi}{2} - a_s)} \tag{1}$$

$$I_{dif,p} = I_{dif,h} \cdot \frac{1 + cos(t)}{2} + aI_{dir,h} + I_{dif,h} \cdot \frac{1 - cos(t)}{2} \tag{2}$$

$$\alpha = arccos(sin(h) \cdot cos(t)\ cos(h) \cdot sin(t) + cos(a_p - a_s)) \tag{3}$$

The nomenclature for Equation 1, Equation 2 and Equation 3 is as following:

- $I_{dir,p}$ is the direct irradiance on the panel.
- $I_{dir,h}$ is the direct irradiance from the sun.

- $I_{dif,p}$ is the diffused irradiance on the panel.
- $I_{dif,h}$ is the diffused ambient irradiance.
- α is the angle of incident between the sun's irradiance and the panel, as denoted by Equation 3
- a is the albedo of the surrounding surface, which denotes how much irradiance is reflected by nearby surfaces.
- a_s is the sun's azimuth, meaning the angle of the sun relative to south (in the northern hemisphere).
- a_p is the panel's azimuth or the angle relative to south.
- t is the panel tilt.
- h is the sun's angular height.

From the panel irradiances and the ambient temperatures, the PV-power is then calculated using the the PV model found in [120]. This model uses a fitted efficiency curve gathered from experimental values, which varies with irradiance and temperature. The curve can be expressed as in Equation 4, and from that the panels output power can be calculated by using Equation 5. In Equation 6 the module's temperature can be estimates from the ambient temperature and the heat by irradiance, if the ambient temperature is assumed to be steady.

$$\eta_{rel}(G', T') = 1 + k_1 ln(G') + k_2(ln(G'))^2 + T'(k_3 + k_4 ln(G') + k_5(ln(G'))^2) + K_6 T'^2 \tag{4}$$

$$P(G, T_{mod}) = P_{STC} \cdot \frac{G}{G_{STC}} \cdot \eta_{rel}(G', T') \tag{5}$$

$$T_{mod} \approx T_{amb} + C_T G \tag{6}$$

Where

- η_{rel} is the real efficiency of the panel at a given temperature and irradiance.
- G is the panel irradiance.
- P_{STC} is the power output at standard test conditions.
- G_{STC} is the irradiance under standard test conditions, usually 1000 W/m^2.
- T_{STC} is the modules temperature under standard conditions, usually 25 °C.
- G' is the normalized irradiance, $\frac{G}{G_{STC}}$.
- T' is the normalized temperature, $T_{mod} - T_{STC}$.

- k_1 to k_6 are curve coefficients found by fitting the model to experimental values.
- C_T is a coefficient determining the heat absorbed by the module from the incomming irradiance.
- T_{amb} is the ambient temperature.

The full PV-module modelling can be found in [120].

For the simulated trade node network there will be five producing nodes creating solar certificates. These simulated nodes are placed in Spain, and are distributed around the country, with a slight bias towards the southern part where there tends to be a higher average irradiance. [121] The simulated nodes are placed on the coordinates given in Table 1, or as seen on the map in Figure 11.

Node nr.	Latitude (North)	Longitude (East)
1	37.3415	-4.3477
2	38.9043	-1.6416
3	39.1649	-5.7183
4	40.6296	-4.4540
5	41.2726	-1.5642

Table 1: The location of the simulated solar energy nodes.

4.1.2 Wind Power Simulations

The trade node network also includes dedicated wind power nodes. These nodes are set to be placed in Germany, and are also simulated using the Renewable Ninja tool. The dataset used for historical wind values is again the MERRA-2 dataset. It provides a relatively fine spatial grid of measurements, and has wind measurements at 2, 10 and 50 meters above the ground. Since wind turbines vary in height, it is necessary to interpolate wind values depending on the height of the turbine. Additionally, a spatial interpolation is done to estimate the appropriate wind values at specific geographical points between grid lines. [122]

The wind values at a given turbine height is interpolated using a logarithmic wind profile, which is practical for estimations sufficiently above ground. The log profile assumes that the turbine is within the atmospheric boundary layer, which is up to roughly 1000 meters above sea level, and that the air has neutral stability. [123] The equation used to create the log wind model can be seen in Equation 7.

$$w(h) = \frac{u*}{k} \cdot \log(\frac{h-d}{z}) = A \cdot \log(h-d) - A \cdot log(z) \tag{7}$$

The wind speed is an equation of the following, as given by [123].

- h is the height above ground.
- u* is the friction velocity.
- k is the Von Karman constant, k = 0.4.

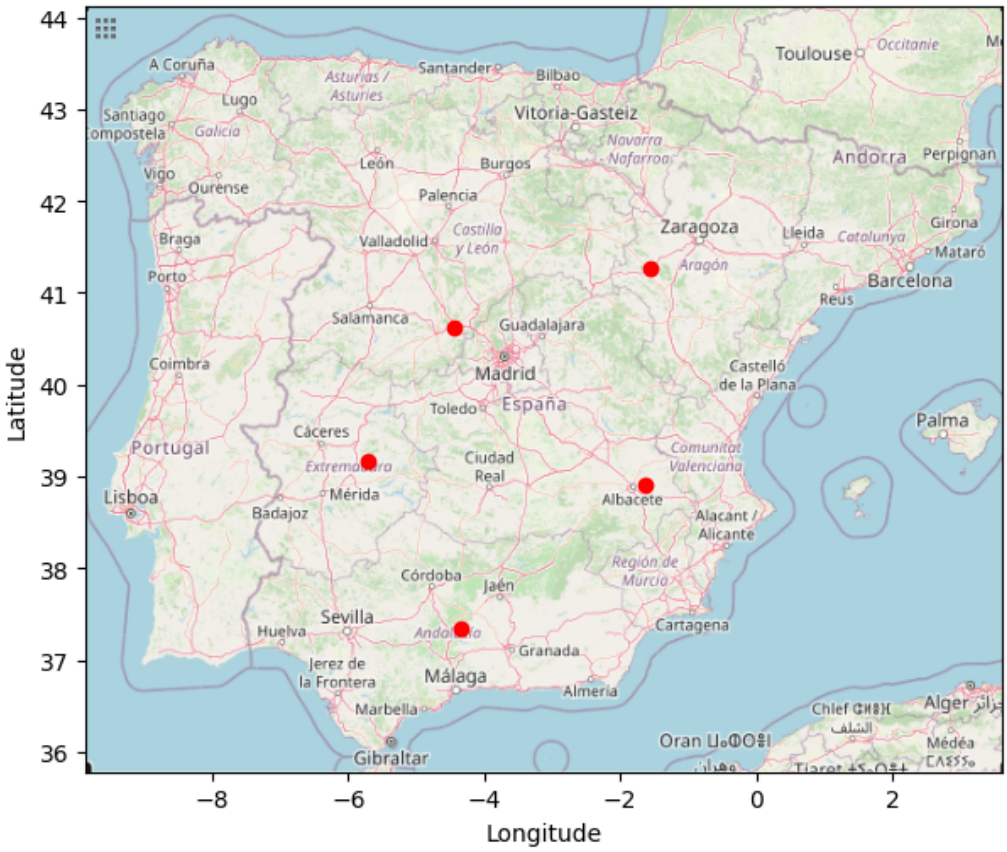

Figure 11: The location of the solar nodes.

- A is the wind sheer, $A = \frac{u*}{k}$.
- d is the displacement height found in the MERRA-2 dataset.
- z is the surface roughness.

In Equation 7 there are two unknowns, A and z. These unknowns can be found found by solving the set of equations created by evaluating Equation 7 at the known points of $h = 2 + d$ meters, $h = 10 + d$ meters and $h = 50$ meters, where d is the displacement height given by the MERRA-2 dataset. The displacement height is meant to adjust the log wind curve such that it accounts for unevenness on the surface level. More details concerning the displacement height can be found in the MERRA-2 FAQ. [124] An example of a log wind curve can be seen in Figure 12, with the MERRA-2 measured points attached. The example log wind curve is for the German wind area 1, located at $(53.3752, 12.9723)$, on the 01.03.2023 from 00:00 to 01:00.

In addition to adjusting the wind estimates for the height of the wind wind turbine, there is also a need to interpolate the wind values between the geographically spaced measuring points provided by MERRA-2. By default, the MERRA-2 measurements are spaced out by latitudinal degrees of 0.5, and longitudinal degrees of 0.625 degrees. The Renewable Ninja tool utilizes a varying Gaussian filter to blur out the grid lines of between measurements. On a two-dimensional scale, a Gaussian filter is applied as a convolution between two matrices, where one of the matrices consists of values taken from a two-dimensional Gaussian distribution, and

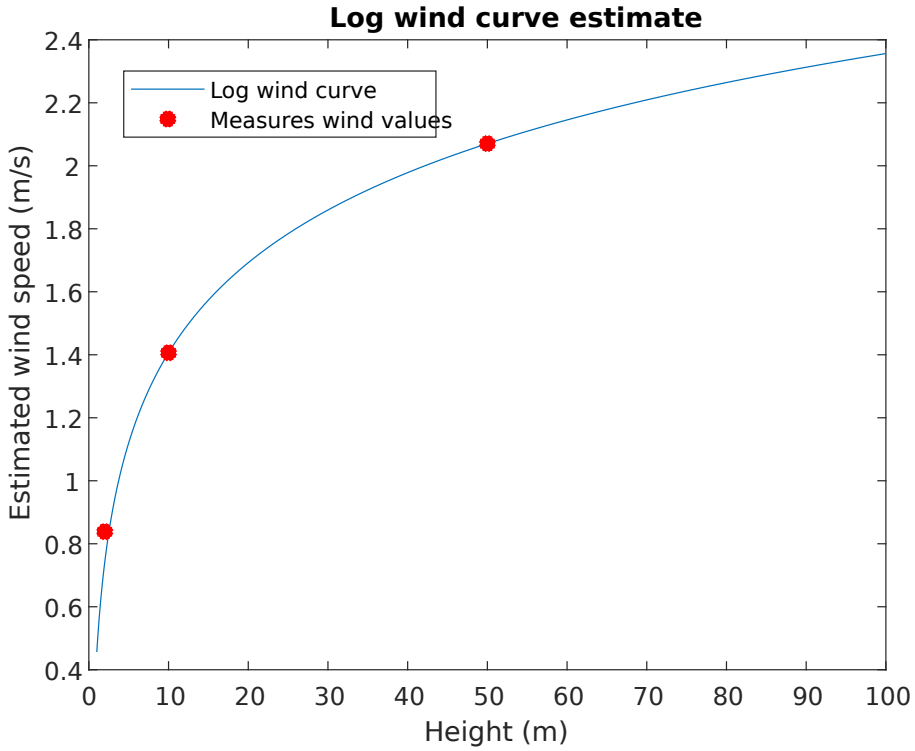

Figure 12: An example log wind curve estimated from MERRA-2 data at wind node nr. 1.

the other matrix is the one consisting of values needed to be blurred. The Gaussian filter used by Renewable Ninja is dependent on the wind speed measured, such that each measurement is the base of a two-dimensional Gaussian distribution centered on the measuring point, and with a standard deviation (filter width) of $\sigma = 0.6 + 0.2v$, where v is the wind speed. A comparative figure is added in Figure 13. This figure consists of three wind value resolutions. The coarsest one is the values as they are given by MERRA-2, the secondly figure shows the measurements with a Gaussian filter applied. For the sake of the demonstration, the filter has a constant filter width of $\sigma = 0.8$. The third figure shows the filtered values at a higher, interpolated resolution, achieving the smoothest wind measurement transitions.

Once good wind values have been calculated for the needed height and geographical location, Renewable Ninja applies a spatial bias. The MERRA-2 dataset is shown to have severe variances between the reported wind measurements, and the locally reported values. However, importantly the general trends and flows are well represented. This means that measurement values can be corrected by applying a geographically dependent correction factor. This applied factor varies from 0.66 to 1.54 within Europe. [122]

Finally, after sufficiently good wind values are obtained, the power outputs can be estimated. For a single turbine, the power output can simply gathered from a power-wind graph provided by the manufacturer of the relevant turbine, or a curve can be estimated from power-wind measurements. In Figure 14, an example of such an estimated curve is shown. The curve is based on the 660 kW turbine by the Danish company Vestas. [126] Next to the estimated curve, is a filtered and normalized version for use in wind farms consisting of this type of turbine. The filtered version is needed to account for internal wind differences in the wind farm, as well as

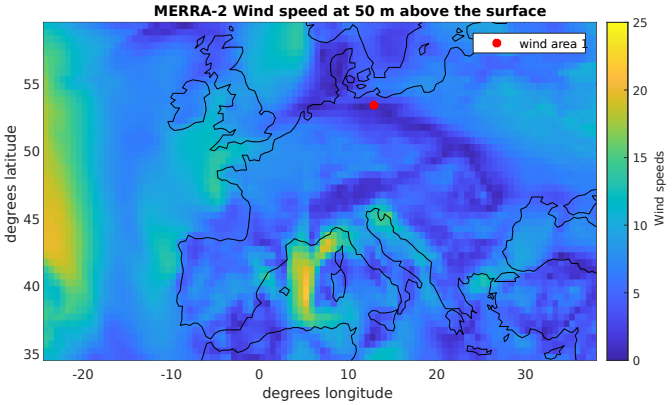

(a) Wind measurements as taken from MERRA-2.

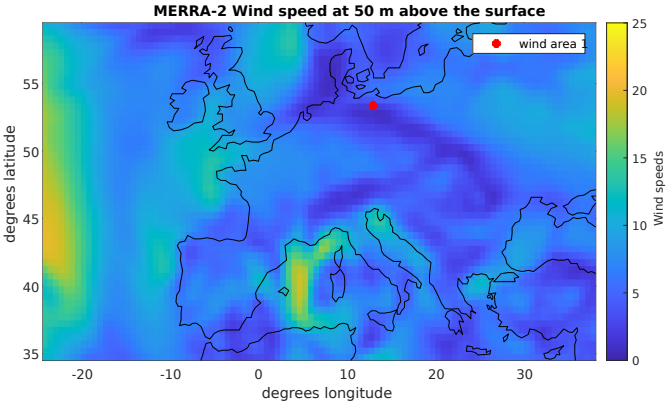

(b) Wind measurements after a Gaussian filter is applied.

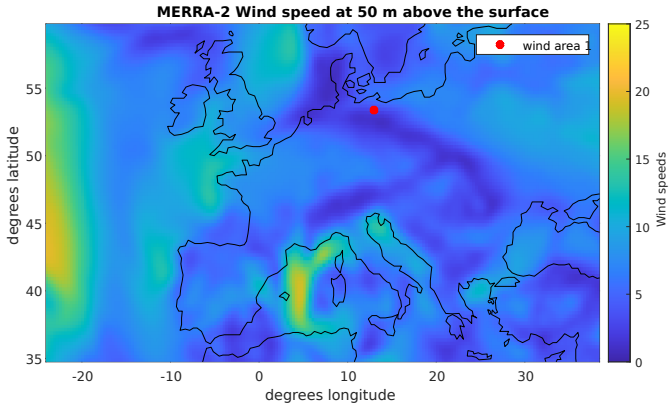

(c) Wind values with both a Gaussian filter and an improved and interpolated resolution.

Figure 13: A comparison between wind values when smoothed and raw. The data is taken from [125].

other varying factors of the turbines themselves.

For the sake of this thesis, the entirety of Germany is simulated as five large wind farms. These simulated wind farms are considered the "German nodes" for certificate trade purposes, and provides the basis for the number of certificates issued on the on-chain marketplace, as well as the time variance the issuing needs to adhere to. The nodes are simulated to be located at the coordinates given in Table 2, or as shown in the map in Figure 15.

4.1.3 Hydro Power Simulations

The third type of simulated certificate injector is hydropower based. With a foundation in the Norwegian power production, five nodes are assumed to produce green hydropower and thereby certificates. Unlike the nodes situated in Spain and Germany, this power production simulation is not done with Renewable Ninja, as it does not offer support for hydropower simulations. Instead, five distinct year-long production time series are drawn from the total Norwegian houerly production for 2022. The hourly production data can be found in [127]. Note Norway's produced energy mix is not completely hydropower, but roughly 89% hydro-based, 10% wind-based and 1% others. [128] Still, as the purpose of these simulations is not

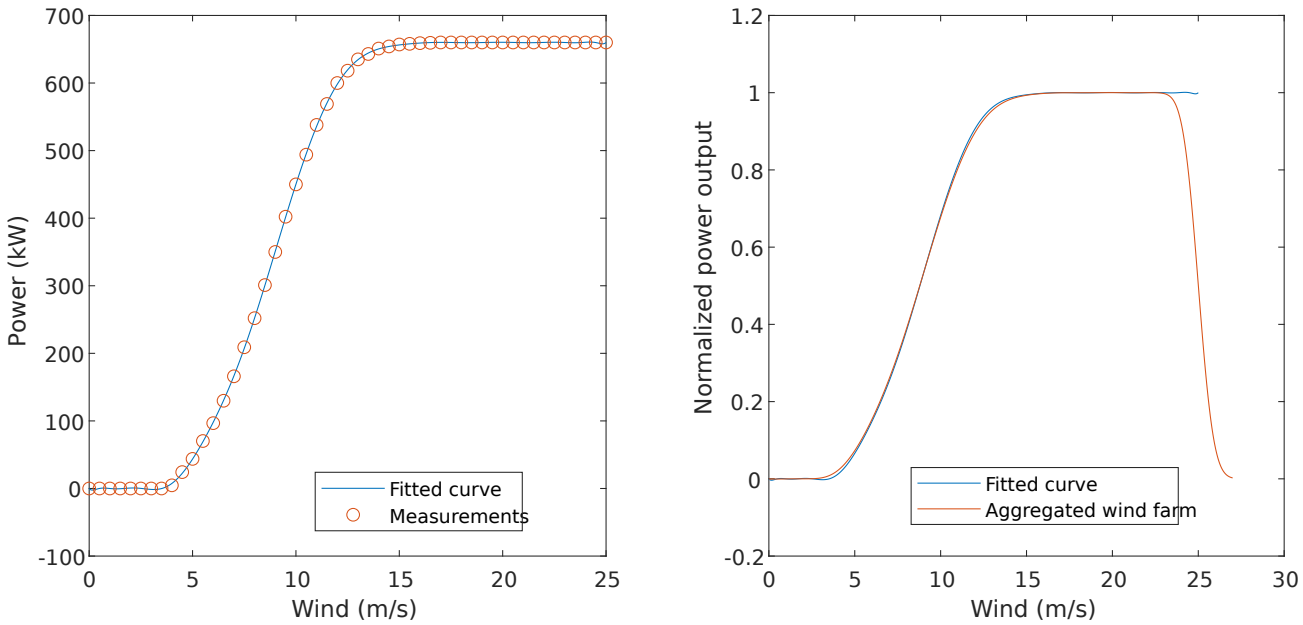

Figure 14: The left figure shows the extracted curve from power measurements for the Vestas 660 kW turbine. [126] The right graph shows a filtered curve that can be used to simulate a wind farm of that turbine type. The graph is filtered by applying a Gaussian filter with a $\sigma = 2$.

Node nr.	Latitude (North)	Longitude (East)
1	53.3752	12.9723
2	51.4703	11.9540
3	52.6793	9.6166
4	50.9776	7.8578
5	50.1544	10.2415

Table 2: The location of the simulated wind energy nodes.

to completely and perfectly emulate power production, but to simply create foundational data for the on-chain marketplace, it is sufficient to assume that certificates from Norway originates from hydropower.

The production of each node in each timestep is found by creating a normal distribution situated around one fifth of the total national production, and then randomly drawing one value according to the created normal distribution. Each normal distribution has an applied standard deviation of $\sigma = 0.03$, such that each node produces unique and distinct, but reasonable values. This leads to power production intervals for each node where two standard deviations, or 95% of values, falls within the range of $[0.14P_{nat}, 0.26P_{nat}]$, where P_{nat} is the national power production at that timestep. As such, the simulated node production at timestep t is given by Equation 8.

$$P^t_{node,hydro} = 0.2P_{nat} \cdot X, \text{ where } X \sim \mathcal{N}(0.2, 0.03^2) \tag{8}$$

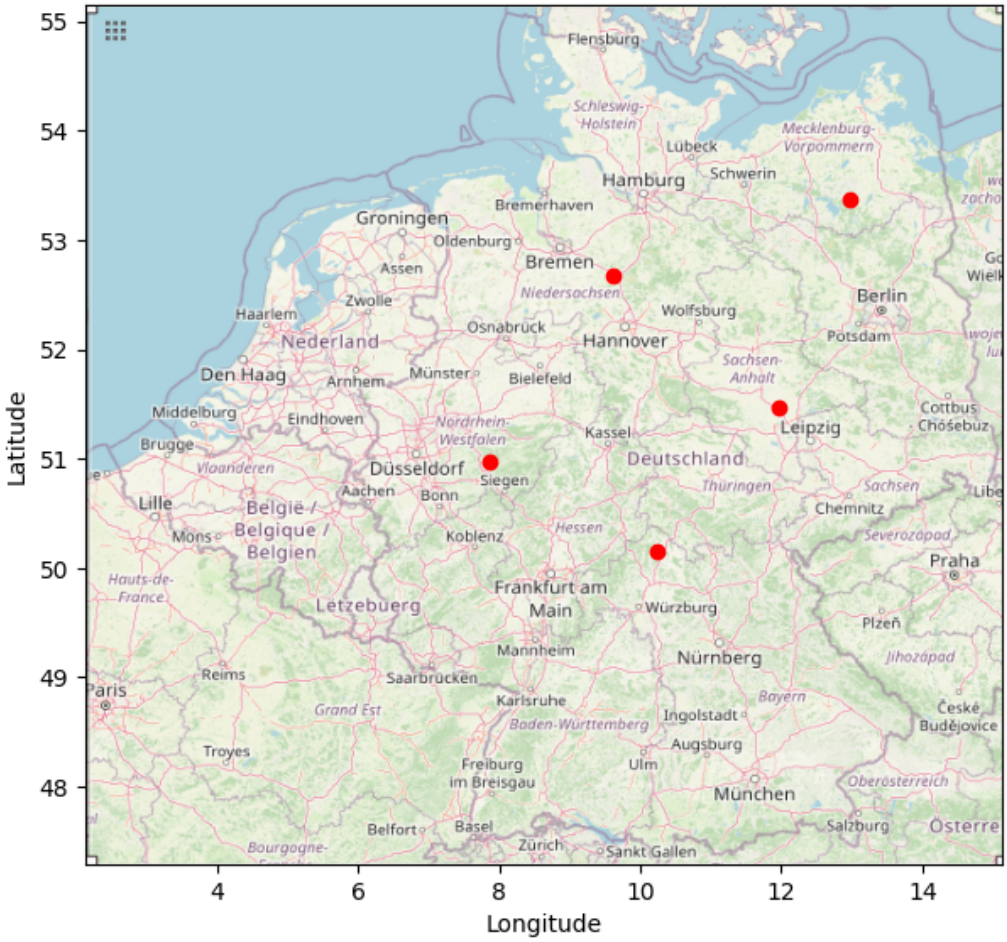

Figure 15: The location of the wind nodes.

4.1.4 Energy Consumption Simulation

Each of the fifteen nodes spread across Norway, Germany and Spain will also act as certificate consumers, such that the there is a steady and realistic inflow and outflow of certificates to the marketplace. The energy consumption of each node is determined by the formula shown in Equation 9. The formula is used to simulate consumption values for a year, and can be broken down into two parts, and two time segments. If the hour $t = 0$, then a value X is drawn from a normal distribution situated around a randomly drawn number within the range [800, 3000], and then multiplied by a seasonal adjustment function $S(t)$, which is described in Equation 10. If the hour $t \neq 0$, then the drawn value X is situated around the consumption value at time $t = t - 1$, and then multiplied with the seasonal adjustment function $S(t)$.

$$C^t_{node} = X \cdot S(t), \text{ where } \begin{cases} X \sim N(\mu, 5^2) & \text{if } t = 0 \\ X \sim N(C^{t-1}_{node}, 5^2) & \text{if } t \neq 0 \end{cases} \tag{9}$$

where

- μ is a random number in the range [800, 3000]

- t is the time in hours

and $S(t)$ is given by:

$$S(t) = 1 + cos^2(\frac{t \cdot \pi}{365 \cdot 24}) \tag{10}$$

The seasonal adjustment function $S(t)$ is needed to adjust for consumption changes typically incurred by winter and summer variations. For better simulations, the seasonal adjustments should differ by country and environment, but for the sake of this thesis, it is sufficient to generalize them for all three node countries. The seasonal variance is adjusted for increased energy consumption in the winter months, as is typical for Norway and Germany [127] [129], which are generally colder countries, but not necessarily the case for Spain, where energy consumption tends to increase in the warmer months.

4.2 The DLT Marketplace

The cyber layer of the complete simulation consists of a DLT marketplace using a blockchain as the common ledger for all involved parties. The goal of the blockchain ledger is to enable a decentralized, but synchronous point which all actors can interact with, eliminating the need for individual peer-to-peer deals and third party interference. Every relevant party to the creation and transaction of a certificate is represented on the marketplace as a wallet address, such that every entity is connected to the blockchain based ledger, and transfers of ownership can be done by altering the blockchain state. This means that in addition to being a secure trading platform, the marketplace can take on the role currently possessed by the AIB trading hub and the Epexspot auction as a hub for both trading and pricing of certificates. By using smart contracts, the GOs can be priced live through the use of proper auction mechanisms, and bids and asks can be automated. For the sake of this thesis, optimal auction mechanisms have not been considered, but for a potential implementation, the role of the auction must be carefully considered and measured. Still, the implemented marketplace is implemented in a fashion similar to common stock trading marketplaces, with a continuous double auction.

Instead the emphasis of the marketplace is to showcase that the auction can be integrated into the blockchain, allowing for near real-time trading which is synchronized with the literal transfer of assets.

The marketplace itself is designed as an interface between the data stored on the blockchain and a client side front end. It is designed with designated roles, delegating each partaker individual rights and abilities. There is a contract owner address acting with the same privileges as the AIB would have, issuer addresses which corresponds to domain issuers, and producer and consumer nodes.

The highest access level is reserved for the AIB, or the contract owner. As they are the central body that oversees the entirety of the EECS they have privileges not held by any other participant. Mainly, the AIB has the privilege to handle and dedicate issuer rights, but it also includes the privilege to execute final certificate trades, after buy- and sell-orders have been matched. Three AIB specific functions are made. The "setIssuer" function, which grants issuer

privileges to a wallet address, and a "removeIssuer" function, which removes issuer privileges, and the "executeTrades" function which executes the final transfer of payment and certificates.

Once issuers have been created by the AIB, the addresses connected to the issuers will have permissions to interact with the issuer level privileges. For the created contract, this is limited to being responsible for issuing and delegating certificates through the "issueCertificates" function. Although, real world privileges also includes pre-approving energy producers as legitimate clean energy producers.

Certificate holder functions apply to all entities handling certificates. This includes producers, but also pure consumers and traders. These functions represents the actions that can be done to already distributed certificates. The "createSellOrder" function adds certificates to a list of purchasable certificates, similarly the "createBuyOrder" function creates an order with pre-approved funds for the purchase of certificates. The "cancelBuyOrders" removes a wallet address' buyOrders from the queue and releases the locked funds, and the "consumeCertificates" function burns the certificates from the blockchain, such that it can be recorded during an off-chain emission accounting.

Finally, an array of functions is available for all users of the smart contracts. These functions are pure information gathering functions, and does not alter the state of the blockchain.

All these functions are implemented using two different standards for NFT. These standards are the ERC721 standard and the ERC1155 standard, respectively used for traditional NFTs and batched NFTs/SFTs [77] [78]. As will be apparent in section 5, the reason for exploring two different standards was that one encountered apparent performance issues for large scale use.

4.2.1 ERC721 Smart Contract

As briefly mentioned in subsection 3.3, there exists an established standard for NFT-handling. This standard is the ERC721 standard [77], which is sanctioned and handled by the official Ethereum organization. The ERC721 standard can be used to implement, issue (mint), handle and transfer NFTs in the typical one token-one tokenId fashion. For the sake of this thesis, the ERC721 standard is one of two standards explored for GO-certificate storage and handling on the blockchain. The fundamental idea of using the ERC721 standard is that each certificate can be uniquely identified, and therefore uniquely tracked as it travels through the network of producers, traders and consumers. With the ERC721 standard, the goal is to create a transparent system where auditing is easy, and malicious certificate accounting is difficult.

Since every token is unique, and is handled by a unique tokenId, all operations needs to be done on every certificate uniquely. This includes amongst others minting, transferring and consuming the certificates. For example, to issue certificates, the blockchain must store a token counter that tracks the number of previously issued certificates, such that newly issued certificates do not conflict with previously issued certificates. Similarly, when a certificate owner wishes to sell tokens, the certificate owner must know, and specify the exact tokenIds to be sold, and subsequently track the changes in owned tokens. Note that this tracking could also be implemented on-chain, but that would add unnecessary gas costs and storage space usage to the contract.

Similarly, to transfer a batch of n certificates, there must be altered an NFT-owner address state on the blockhain for each certificate transferred. As will be apparent in section 5, this is transactionally heavy and space complex. Figure 16 shows the flow of certificates from a producer owning certificates, to a consumer buying certificates. Note how the certificates have different tokenIds, and requires n transactions for n certificates.

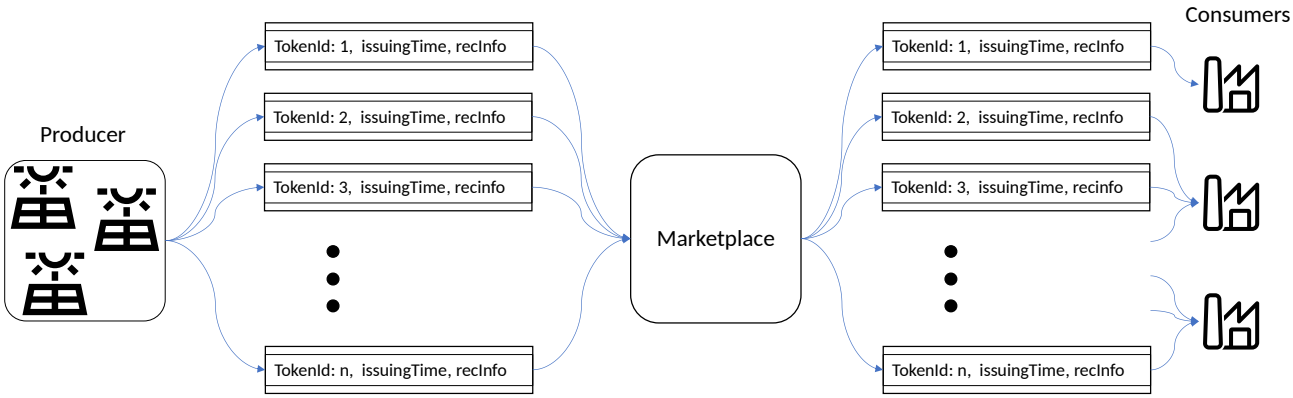

Figure 16: An example flow of ERC721-based certificates from producer to consumer.

4.2.2 ERC1155 Smart Contract

The alternative standard implemented is the ERC1155 standard. This is a newer standard, and is the standard for SFTs. [130] The ERC1155 standard implements tokens that fall somewhere between traditional fungible tokens, such as used in the crypto currency space, and the non-fungible tokens used under the ERC721 standard. The SFT space consists of tokens marked at a batch by batch level, such that groups of tokens are uniquely distinct from other groups of tokens, but tokens within a group are indistinguishable and interchangeable. This means that some information is lost with regards to tracking and transparency. Unlike pure NFTs, it is possible to scramble certificate trading paths if the certificates originated from the same pool. In this case by coming from energy produced in the same place, during the same hour, and from the same energy source. I.e hydro, wind or solar.

Since certificates are able to be batched together with this standard, handling can be done at a batch level, meaning that multiple certificates can be issued, transferred or consumed in a single transaction. This significantly reduces the overhead required by the smart contract and user, as address-tokenId relations does not have to be maintained in the same manner as with ERC721 tokens. Furthermore, sell orders and buy orders can be simplified into a single order type, as they now both can manage a general amount of tokens, rather than a complete list of IDs. This is tried visualized in Figure 17, where there only needs to be a single transaction to add multiple certificates to the marketplace, and there only needs to be one transaction for each outgoing order as well, rather than one for each certificate.

4.2.3 Deploying, Connecting and Interacting with the Smart Contract

The smart contract is situated on the blockchain, within the EVM, and needs to be connected to such that it contact functions and features can be accessed on local computers, and actions can be automated. The library "Web3Py" is a Python library designed for this purpose. [131]

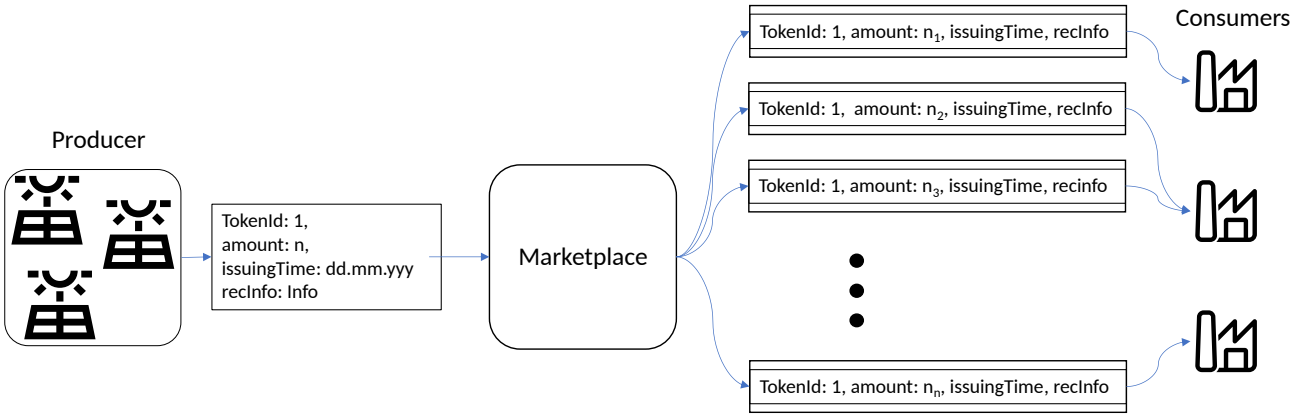

Figure 17: An example flow of ERC1155-based certificates from producer to consumer.

The use of the Web3py library is two-fold. There needs to be a contract object that can be interacted with, and then proper transaction calls need to be made to the contract instance to initiate state changes on the blockchain.

For this thesis, the blockchain was simulated using the open source tool Ganache-cli. [132] Ganache-cli is a NodeJs based blockchain simulation that mimics the behavior of a private blockchain. Ganache-cli is command line based, and offers a wide variety of customizable options, allowing the user access to for example pre-populated test wallets that can be written to a local .json file. This is a very useful feature when simulating multiple nodes with corresponding wallet addresses. An example of the most used Ganache-cli settings for this thesis can be found in Listing 1.

```
ganache-cli --logging.debug=true --logging.quiet=false --accounts 100 --
    defaultBalanceEther 10000 --acctKeys ".\ganacheAccountsERC1155.json"
    --port 7788 --gasLimit 2000000000 --gasPrice 200000000 --database.
    dbPath "./chainDatabase"  --mnemonic
```

Listing 1: anache-cli batch command to run the local blockchain simulation with typical settings.

Before the contract can be deployed onto the blockchain, it needs to be compiled down into its opcode and bytecode form. The opcodes and bytecodes are respective machine instructions for to be run on the EVM, where the opcodes are human readable low level machine instructions, and bytecodes are hexadecimal representations of that same code. [133] The Python library "py-solc-x" includes a Python friendly port of the solc Solidity compiler that can be specified for a given solidity version. [134]. Solidity is a language under development, and due to rapid changes the compiler used is version specific. In addition to the opcode and bytecode representation of the contract, the compiled contract also includes the Application Binary Interface (ABI), which is an interface for the functions and types available in the contract.

The Web3Py library connects to the blockchain through an HTTP connection. Since the Ganache-cli local blockchain is situated on the computer, the connection point is the localhost IP: 127.0.0.1, and the connection port is the specified port 7788. Each chain must also be specified with a unique chain ID, which in this case is set to be 1337. Once there has been established a connection with the chain, the compiled contract can be deployed by creating a

contract object from the ABI and bytecodes from the compiled contract, and then creating a transaction that calls the constructor method within the contract. See Listing 2 for a code example of a transaction calling the constructor method.

```
    # initialize contract
    RECMarket = w3.eth.contract(abi=abi, bytecode=bytecode)
    nonce = w3.eth.get_transaction_count(ownerAddress)
    # set up transaction from constructor which executes when firstly

    transaction = RECMarket.constructor().build_transaction({
        "gas": w3.eth.gas_price,
        'maxFeePerGas': maxFeesPerGas,
        "chainId": chain_id,
        "from": ownerAddress,
        "nonce": nonce,
    })

    signed_tx = w3.eth.account.sign_transaction(transaction, private_key=
       ownerPrivKey)
    tx_hash = w3.eth.send_raw_transaction(signed_tx.rawTransaction)
    tx_receipt = w3.eth.wait_for_transaction_receipt(tx_hash)
```

Listing 2: A raw trasaction formed using the web3Py library. It calls the constructor of the contract, placing the contract on the chain.

The transaction in Listing 2 is built as a raw transaction. Within the transaction there are some keywords that have been specified. The transaction needs a gas price, which is the cost of each gas used by the transaction, a "maxFeePerGas", which must be less than the gas price, a chain ID that specifies which chain the transaction is executed on, and a nonce. The nonce here is not the same as the nonce often associated with a PoW chain, where it is important in proving the validity of some done work. Here the nonce is equal to the number of transactions done by the transacting wallet address, and is important in order to prevent double spending, and correctly order transactions occurring within the same block from the same address. [135] After the transaction is built, it is signed and sent to the pool of transactions before being approved.

Just as the constructor method of the contract can be called through the use of a raw transaction, the contract functions can be extracted as well. To be able to automate on-chain transactions through python calls, a library of python functions creating and calling raw transactions was made. To set up a library that is able to create transactions and call smart contract functions, it is necessary to obtain the contract object situated on the blockchain. In Listing 3, the code to get the contract object is shown. The contract object ("DoubleAuctionForRECs"), has two main callable objects: 1) functions and 2) events. If the functions alter the state of the blockchain, then the functions are called through the already seen raw transactions, otherwise they can be called with the simple .view() method. Events are signals that are emitted to the function caller, and is a way to return information from the blockchain. Events are typically used to confirm that a transaction has occurred, and to return some useful information about the new state of the blockchain. With the contract object, smart contract functions can be called in a similar manner as with the constructor.

```

```

```
from web3 import Web3 #Had to install web3 directly from github, as the
    current stable version is depricated
import json

w3 = Web3(Web3.HTTPProvider("HTTP://127.0.0.1:7788", request_kwargs={'
    timeout': 2000}))
chain_id = 1337
maxFeePerGas = 2000000000
#get abi
with open("abiERC1155.json", "r") as file:
    abi = json.load(file)

#get contract info
with open("contractInfoERC1155.json", "r") as file:
    contractInfo = json.load(file)

DoubleAucitonForRECS = w3.eth.contract(abi = abi, address = contractInfo["
    contractAddress"])
```

Listing 3: A code snippet showing how to access the smart contract object existing on the blockchain.

4.2.4 Order of Operations for the Node Network Simulations

The node network is simulated on an hour by hour basis, with the production and consumption data of each node as the inputs of the simulation. The simulaiton is set up in a synchronous matter, such that one operations is allways followed by the same other operation. This works well, but it does inherently create some consistency in the simulation results that would not exists if the different orders came in at sudo-random intervals. The sudo-code presented in Listing 4 shows the general order of operations used for the simulations. For each step there is used an array of other functionalities to capture and store the data at each hourly time-step, and to decide which certificate types should be selected for the current operation. Still, the sudo-code encapsulates the general flow of the simulation.

```
    for each hour in a year:
        for each node in Germany:
            issue certificates to node according to production data
            place sell order for certificates produced
            place buy order for certificates to consume
        for each node in Spain:
            issue certificates to node according to production data
            place sell order for certificates produced
            place buy order for certificates to consume
        for each node in Norway:
            issue certificates to node according to production data
            place sell order for certificates produced
            place buy order for certificates to consume

        match buy and sell orders
        execute trades of certificates

        for each node in all nodes:
            consume certificate according to consumption data
```

Listing 4: The order operations for each node during the network simulations.

4.2.5 Performance Evaluation Metrics

As this thesis proposes two standards for on-chain certificate handling, it is useful to have concrete evaluation criteria that can be used to compare the two standards. Therefore, the two standards are evaluated on the two concrete metrics of storage space complexity per certificate handled, and time complexity per certificate handled. These concrete criteria need to be evaluated against the less concrete, but very important metric of traceability and transparency. These criterias are closely linked to the typical blockchain performance criteria of transactions per second, and total storage space occupied. These metrics are both common bottlenecks for on-chain migration, and should in general be considered when evaluating the efficiency and feasibility of migrating an off-chain service into a blockchain.

The space usage of a transaction can be estimated by from the gas of a transaction. The transaction's gas is a function of a constant transaction gas cost, the variable costs tied to the opcodes executed and the added space taken up on the blockchain. A complete overview of gas use of different operations can be found in Ethereum's yellow paper, however the majority of gas used comes from commonly executed opcodes with gas values grouped into zero-, base-, very low-, low-, mid- or high-cost groups. [109] Each opcode group has different gas values, which are as seen in Table 3. Some opcodes do not fall into these groups, and have custom pricing, most importantly the "CREATE" opcode, which is used when a contract is deployed, and costs 32 000 gas by itself. Additionally, some gas costs are flat fees, such as a condition based fee of 9000 gas when a value transfer is enacted, or a constant fee of 21 000 gas for every initiated transaction. Finally, after all operational fees have been accounted for, the transaction size can be estimated by knowing that each created byte of a new contract costs 200 gas, and that each non-zero byte in a regular transaction costs 16 gas.

Type	**Name**	**Gas value**	**Frequency**
Opcode group	Zero	0	per opcode called
Opcode group	Base	2	per opcode called
Opcode group	Very Low	3	per opcode called
Opcode group	Low	5	per opcode called
Opcode group	Mid	8	per opcode called
Opcode group	High	10	per opcode called
Flat fee	CREATE	32 000	per contract created
Flat fee	SUICIDE	5 000	per opcode called
Flat fee	Call value	9 000	per value transfer
Flat fee	Transaction fee	21 000	per transaction
Size dependent fee	Contract size fee	200	per contract byte
Size dependent fee	Transaction size fee	16	per non-zero transaction byte

Table 3: An overview of the gas costs of opcode groups and significant other costs. [109]

The time efficiencies of each contract is measured and compared as well. This is done by running tests with set amounts of certificates on the same hardware. The tests are run on a personal Acer Aspire S13 laptop from 2016 similar to the one found in [136]. It has an Intel Core i7-6500U 2x2.5 - 3.1 GHz processor and and Intel HD Graphics 520 graphics adapter. By default, the Ganache-cli software executes transactions instantaneously in order for the test environment to be efficient and improve the run time. However, when measuring the time complexity of

each contract the emulated environment should be as similar to a real world implementation as possible, and account for the inter-block time that exists on Ethereum blockchains. This can be done by running Ganache-cli with the "–blocktime" flag to manually set inter-block times. For Ethereum based chains, the inter-block time is 12 seconds.

5 Results

5.1 Certificate Production

In order to feed the node network with certificate production data, respective simulated time series data for each certificate type: wind, hydro and solar, has been produced. The data is in its entirety shown as graphs in appendix A.1, A.2 and A.3, where each plot corresponds to a node/area, and each group of plots correspond to a country. A subset consisting of the wind certificates from area 1 in Germany, the solar certificates from area 1 in Spain, and hydro certificate from area 1 in Norway is presented in Figure 18.

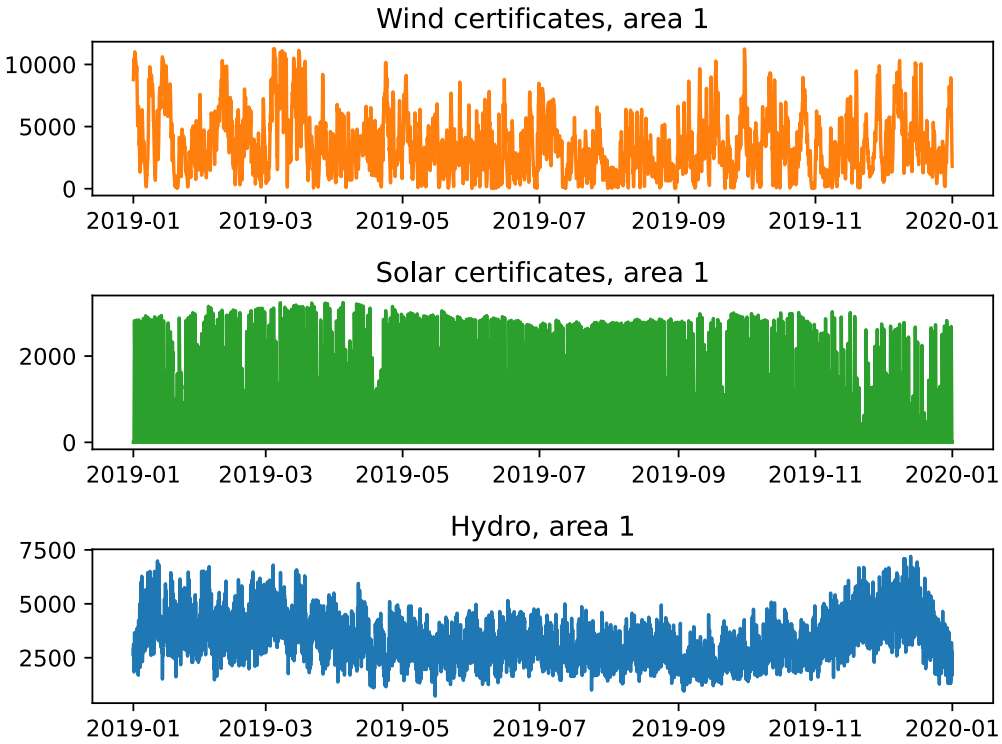

Figure 18: The certificate production profiles of area 1 for the wind, solar and hydro simulations.

The wind profile is sporadic, and rarely hits peak values. Germany's absolute max wind capacity is 58 100 MW, resulting in a node max of roughly 11 620 MW, or an equivalent number of certificates. [137] The solar time series is more periodic, with pretty consistent peaks and valleys. There is a tendency of seasonality, with a dense certificate production in the summer months, and a less dense production in the winter months. The Spanish nodes produce a maximum of roughly 3800 certificates per node per hour. [138] The hydro power time series is semi-sporadic, but follows a clear seasonal pattern with higher production in the winter months. Hydro power is generally considered a flexible energy source that can follow the consumption pattern. However, at the end of the hydro time series there is a quick dip in production. This dip is a result of the overlying production curve for Norway having a quick production dip in December 2022 and January 2023. [127] Comparing this dip to earlier years, it becomes apparent that the dip around late December-early January is consistent and annual, and it is

though that this dip comes from the Norwegian industry being on holidays. [139]

To see and confirm the general behavior of each power source, a 14 day interval of the wind, solar and hydro curves are shown in Figure 19 Note how the wind curve is sporadic and highly varying, the hydro curve is varying, but less than the wind curve, and the solar curve is consistently following day-night cycles.

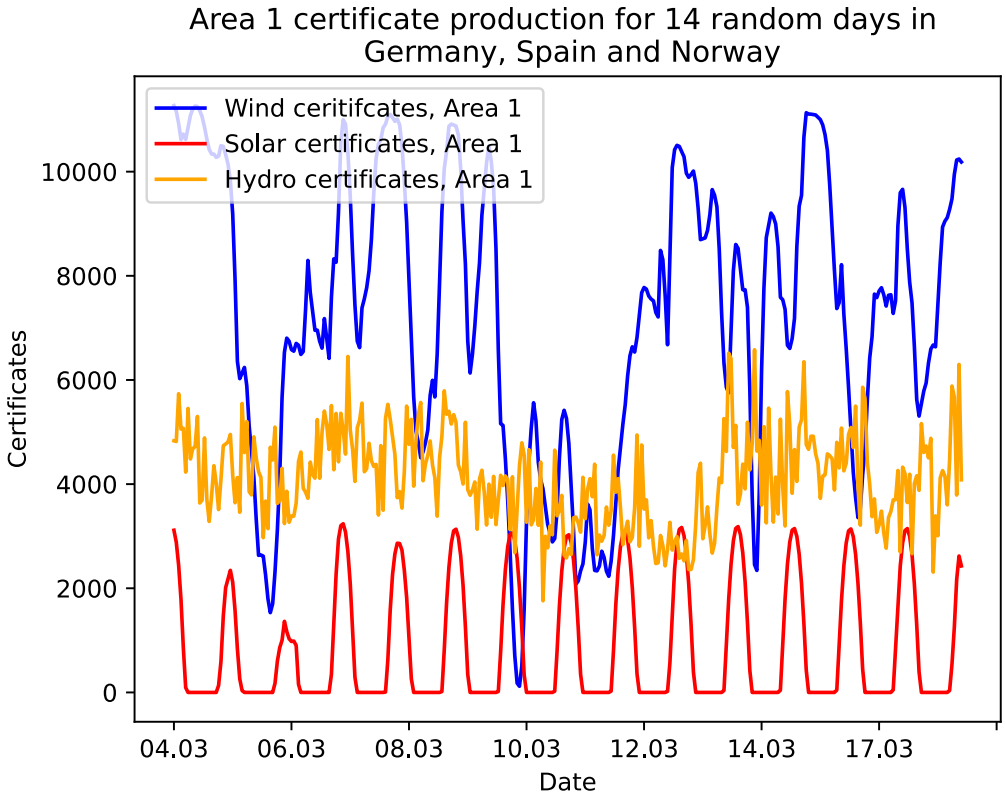

Figure 19: Certificates produced by by solar, wind and hydro over a random 14 day period from the Spanish, German and Norwegian area 1.

5.2 Consumption Data

The consumption data is produced following the method detailed in subsubsection 4.1.4. The data produced consists of fifteen consumption time series, where each follow a seasonal pattern. Otherwise, they have sudo-random patterns, which introduces uncertainty and unpredictability to the system. An example of this is the lowest consumption curve in Figure 20, which develops an uncharacteristic curve compared to the other consumption curves.The consumption data in these time series are used to govern the certificate consumption in the node network.

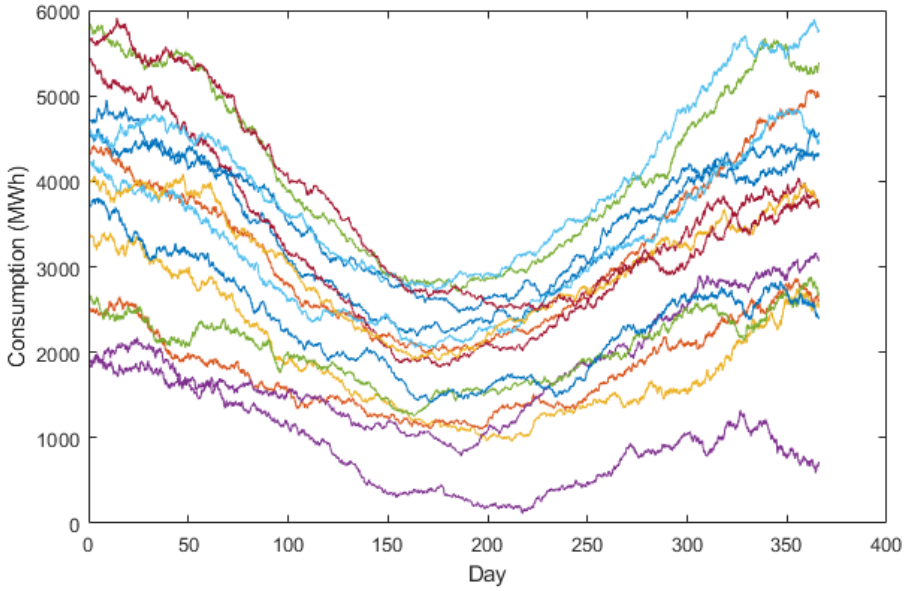

Figure 20: The consumption data used for the 15 nodes.

5.3 ERC721-Based Marketplace

The ERC721 based marketplace built upon the idea that each certificate was to have its own NFT representation on the blockchain. As explained in subsubsection 4.2.1, this means that there are independent operations done on on every certificate for every transaction. This provides traceability, but it also comes with an expected increase in space use and time use. The contract has certificate dependent operations, and certificate independent operations. The two categories are presented in separate tables, measuring values at 1, 10, 100 and 1 000 certificates. These values are far below that used in a real world certificate marketplace, but testing for this standard had to be limited due to both very long run times of large certificate amounts, and fully occupying the memory of the computer.

5.3.1 Storage Space Complexity

To measure the space use, each transaction was placed in a separate block. This means that each measurement is a combination of the space added by the transactions itself, and the inherent space use of a block, such as the nonce, previous hash, difficulty, merkle root and more.

In Table 4, the space use of certificate independent functions are presented. The contract deployment is cost heavy, and at 22.4 kb it is close to the 24 kb constraint set on contracts deployed on the Ethereum network. [140] This means that on the proper Ethereum mainnet, the contract would be expensive to launch, due to high costs. However, on a private and permissioned blockchain, the gas limits, and therefore the size limits are customizable and adaptable. The other fixed space costs are minor, at between 600 and 700 bytes, where the block information itself makes up 509 of of those bytes.

Action	Space costs (bytes)
Empty block	509
Contract deployment	22400
Assigning the issuer role	651
Create a buy order	693
Cancel all buy orders	620

Table 4: ERC721 - Space costs of functions independent on the number of certificates, in bytes.

The certificate dependent space use is shown in Table 5. Surprisingly, the issuing of certificates and execution of trades are relatively unaffected by the number of certificates, and are primarily dominated by the space cost of creating a new block. Comparatively, the space costs of creating sell orders and consuming certificates grow with the number of certificates.

By fitting a curve to these points, it is possible to estimate the size of the blockchain after one year of of trading. Assuming that every certificate created is added to a sell order, and that every MWh of consumed energy corresponds to a consumed certificate, the blockchain size after one year would be roughly 19.67 Gigabytes. This accounts for the total number of certificates produced, which is 314 690 720 for the fifteen nodes spread across Germany, Spain and Norway. This is a significant, but not unmanageable size.

Action	1 certs.	10 certs.	100 certs.	1000 certs
Issuance of certificate	945	945	945	947
Create a sell order [1]	1434	1704	4514	32644
Execute trades	621	621	621	623
Consume certificates	716	1009	3800	31920

Table 5: ERC721 - Space costs of functions dependent the number of certificates, in bytes.

5.3.2 Time Complexity

The runtimes of the different functions are presented below. All runs are done on the same computer, an Acer Aspire S13 laptop from 2016 similar to the one found in [136].

The time complexity of the independent functions are trivial. Since the blockchain operate with 12 second inter-block times, any runtime less than 12 seconds is insignificant and cannot reasonably be measured. This is the case in Table 6, where all the time taken by each operation is dominated by the inter-block time, rather than by the complexity of the operation. In rare cases, it is possible to get time over 12 seconds for these independent functions, if the previous block was deployed just as the operation started. Then there are runtimes of between 12 and 13 seconds. However, these are rare, and overall the time use of certificate independent functions are negligible compared to the inter-block time.

[1] The sell order function is actually comprised of two functions, one that allows the contract to transfer the NFTs on behalf of the owner, and one to create and store the order.

Action	Time (s)
Contract deployment	< 12
Assigning the issuer role	< 12
Create a buy order	< 12
Cancel all buy orders	< 12

Table 6: ERC721 - Time use of functions independent on the number of certificates, in seconds.

The runtimes presented in Table 7 changes from being dominated by the inter-block time for low amounts of certificates, to being dominated by the operation complexity for larger certificate amounts. While operating with less than 10 certificates, the runtime of each operation is less than one inter-block time, and it does not make sense to give any time estimate other than less than 12 seconds. When handling more than 100 certificates, there is a significant reliance on computational power and blockchain transactional speeds, especially with 1000+ certificates. Note that due to the inter-block time, there is a significant variance of up to 12 seconds, depending on when the transaction is added and if it corresponds well with the next block. Furthermore, it clearly scales with the number of certificates for all functions, which differs from the space costs where issuing certificates and executing trades did not take up significantly more space when handling more certificates.

Action	1 certs.	10 certs.	100 certs.	1000 certs.
Issuance of certificate	< 12	< 12	77.3	768.1
Create a sell order	< 12	< 12	29.4	203.8
Execute trades	< 12	< 12	78.9	890.6
Consume certificates	< 12	< 12	17.9	271.4

Table 7: ERC721 - Time use of functions dependent the number of certificates, in seconds. All times are subject to a variance of 12 seconds, due to the inter-block time of the network.

5.4 ERC1155-Based Marketplace

Similarly to the ERC721 testing, each transaction was placed in a separate block, meaning that the space occupied includes that of an empty block, and that transaction times include potential inter-block time. The same measurements of 1, 100, 1 000 and 10 000 certificates are done, but due to the increased speed of this standard, there are also measurements for 100 000 certificates.

5.4.1 Storage Space Complexity

Table 8 shows the certificate independent functions of the ERC1155 standard. Again, the contract deployment occupies a large amount of space, at 25.54 kb. This is more than the Ethereum mainnet limit, but it is deployable on the simulated Ganache cli blockchain, and would be on a private and customizable chain as well. The other functions are slightly more space heavy as well, when compared to their ERC721 counterpart, but overall still dominated by the block information itself.

Action	Space costs (bytes)
Empty block	509
Contract deployment	25540
Assigning the issuer role	651
Create a buy order	793
Cancel all buy orders	717

Table 8: ERC1155 - Space costs of functions independent on the number of certificates, in bytes.

The ERC1155 standard employs batched NFTs, or SFT, and offer a significant space cost reduction when scaling to a large number of certificates. This is confirmed when measuring the space occupied with increasing certificates. Table 9 shows the space cost added to the blockchain with a single transactions handling up to a 100 000 certificates. The results shows that scaling the number of certificates does not increase the space occupied on the blockchain by that single transaction. However repeated transactions will still increase the storage space use.

Action	1 certs.	10 certs.	100 certs.	1000 certs.	10 000 certs.	100 000 certs.
Issuance of certificate	782	782	782	782	782	782
Create a sell order	1468	1468	1468	1468	1468	1468
Execute trades	621	621	621	621	621	621
Consume certificate	747	747	747	747	748	748

Table 9: ERC1155 - Space costs of functions dependent the number of certificates, in bytes.

5.4.2 Time Complexity

The time complexity measurements for the ERC1155 standard is constant, and does also not depend on the number of certificates. All single transactions take less than 12 seconds to run, and are therefore dominated by the inter-block time, rather than the runtime of the transactions themselves. This means that all single transactions, independent of the number of certificates, takes somewhere between 0 to 12 seconds to execute, depending on when the next block is added to the blockchain.

5.5 Node Network Behavior

This section presents and compares the behavior data for the nodes in Area 1 of Spain, Norway and Germany. Similar data exists for the remaining nodes, but is not showcased here due to their very comparable nature. The general trends explained here are consistent across all the similar areas. From the ERC721 and ERC1155 time and space complexity results, it became apparent that the time complexity of the ERC721 standard was too problematic for a large scale simulation. Therefore, the results presented below are for a node network using the ERC1155 standard. The node network has simulated certificate production, trade and consumption for a year, and it took roughly 72 hours for the full simulation.

The marketplace was fed with production and consumption data, and programmed to buy and sell certificates as needed. Each node is set to buy the energy origin certificate that dominates the market. I.e. the certificate type that has the most registered sell orders on the blockchain. Such behavior assumes and encourages frequent and fluid trading, but does not assume that nodes prioritize using self-produced certificates first. Given that there are certificates on the market, the nodes will try to cover their entire energy consumption with certificates. In Figure 21, the certificate consumption is plotted for each type of certificate: wind, solar and hydro. For comparison, the energy consumption of that node is plotted with it.

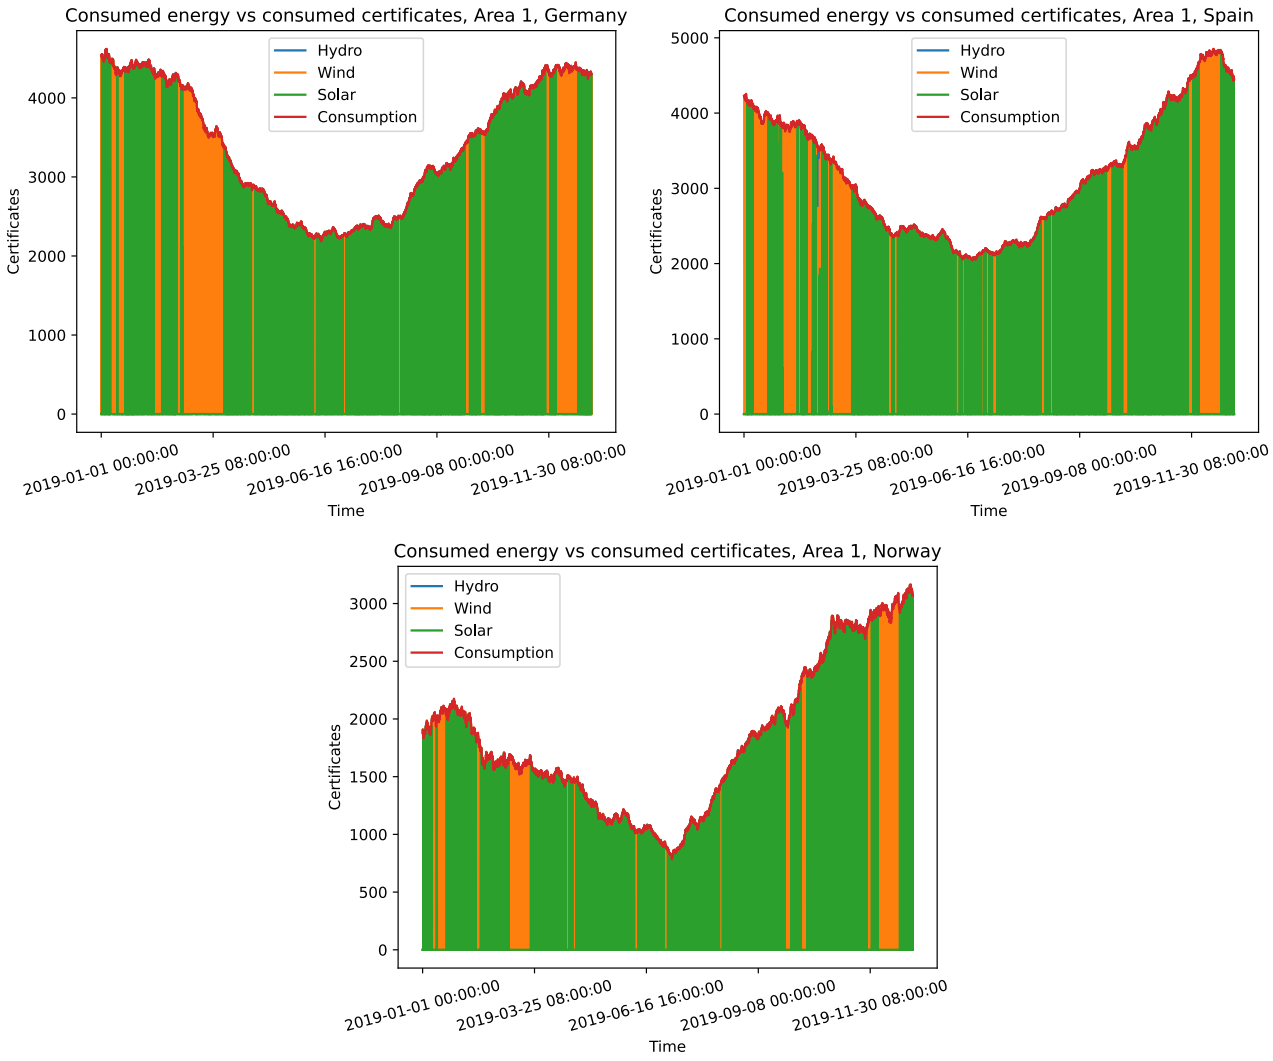

Figure 21: The energy certificates consumed by node 1 in Germany, Spain and Norway. The consumption curve is overlaid to show how the certificate consumption match the energy consumption.

From these plots, it seems that these nodes majorly consumes solar certificates, however that is inaccurate. These plots nicely compares the certificate consumption with the energy consumption, but the certificate consumption plots for each certificate type are overlaying each other. Furthermore, the dense nature of these plot hides areas where there are not enough certificates on the market, and the nodes are limited in their certificate consumption.

A variety of alternative figures are presented to elaborate on the concrete behavior of the marketplace under the simulated conditions.

In Figure 22 a 24 hour interval of certificate consumption is presented for each country. Compared to the plot of the yearly consumption, this shows that there is great variety in the type of certificate consumed. Within this cutout, there tends to be enough certificates on the marketplace to fully cover the energy consumption with a single type of certificate. This indicates a certificate surplus on the marketplace. The 24 hour cutout is from June, where the energy consumption is modeled to be low, so the certificate surplus can be expected.

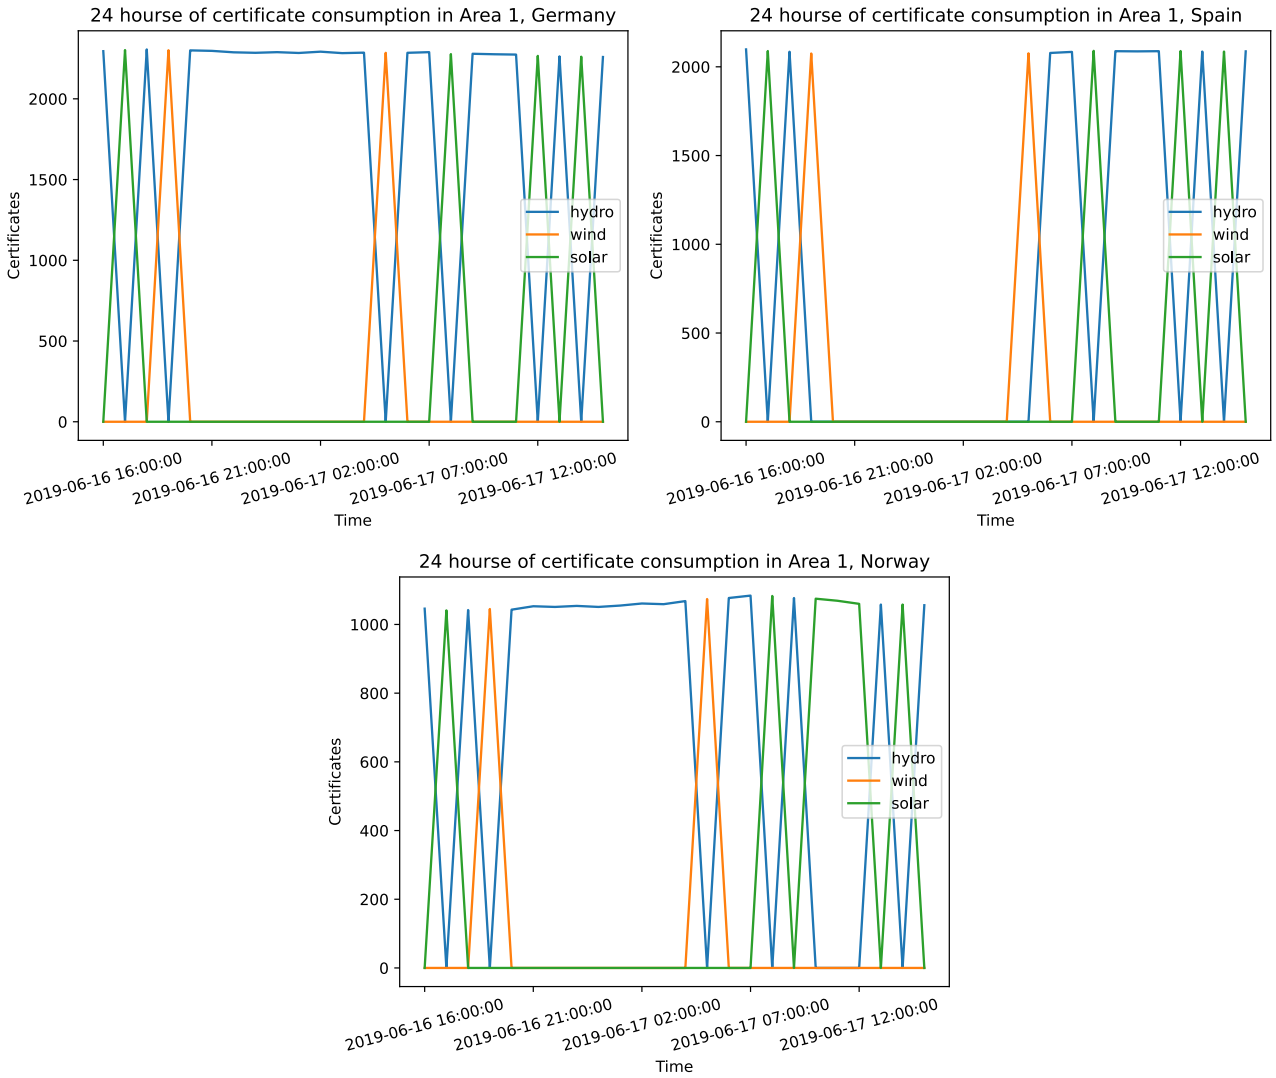

Figure 22: A random 24 hour interval certificate consumption.

In Figure 23 the certificate consumption is plotted as scatter plots instead. From these plots it is more apparent that there are periods with certificate deficits in the first three months of the year. During these deficit months, each node must buy either a mix of certificates to fulfill the consumption quota, or is simply not able to consume enough certificates to cover their energy consumption. Again, note that just as in Figure 21, there is an overlap of plots, and it is difficult to determine which certificate type dominates.

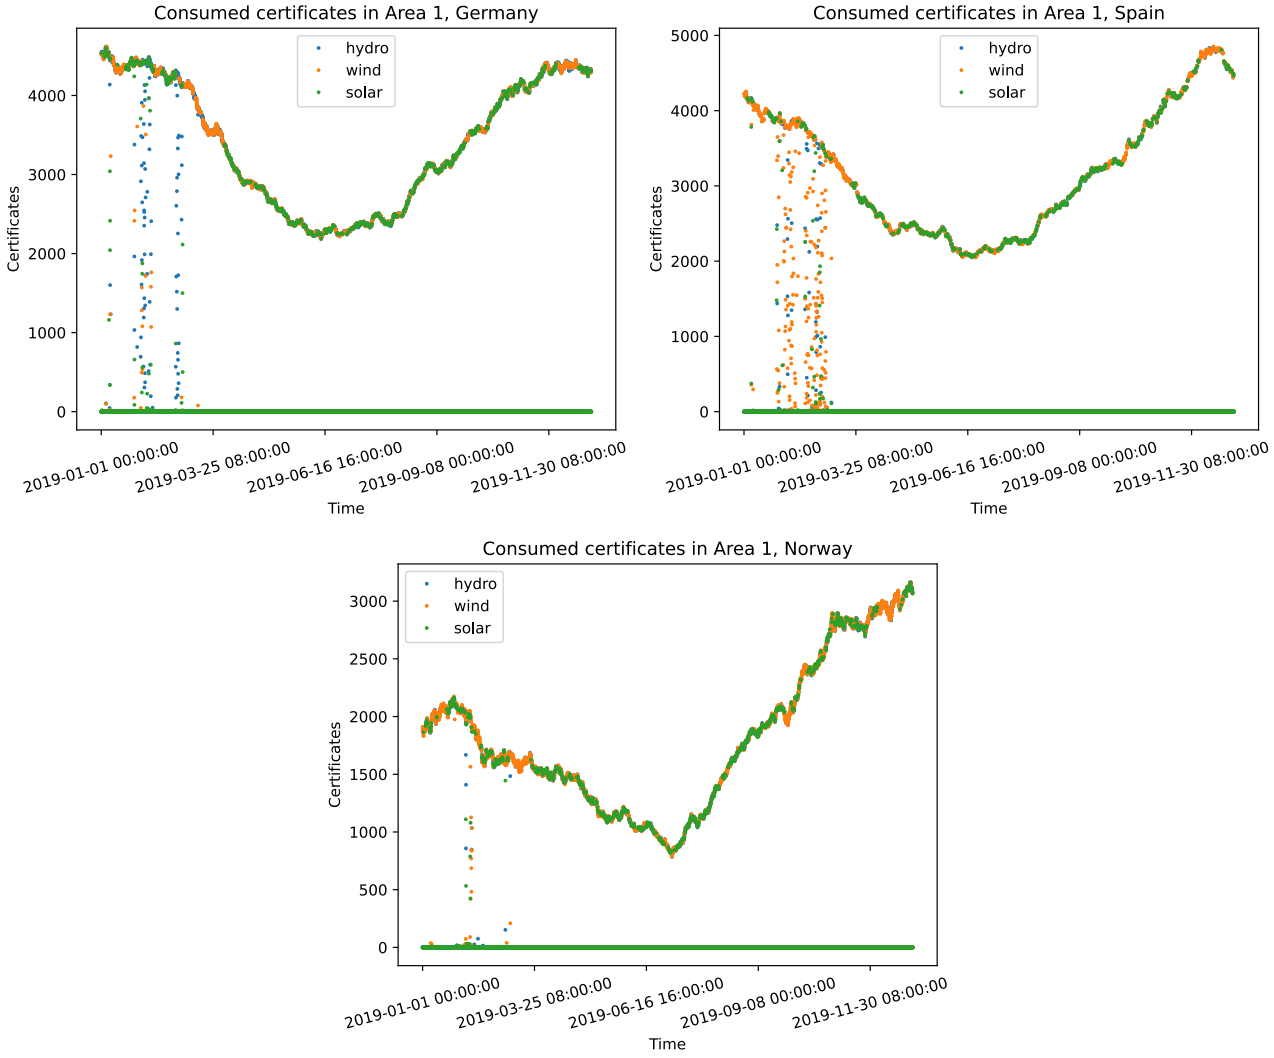

Figure 23: A scatterplot of the certificate consumption, better showing the times where there is a certificate deficit in the market.

As both Figure 21 and Figure 23 have overlaying plots, it is useful to decompose each certificate consumption curve into their respective certificate types. This is done in Figure 24. These graphs show the distribution of certificate consumption better, and confirms that all certificates are traded throughout most of the year. There seems to be slight gaps where the solar certificates are scarce, which is expected as the scarcity happens during the winter months.

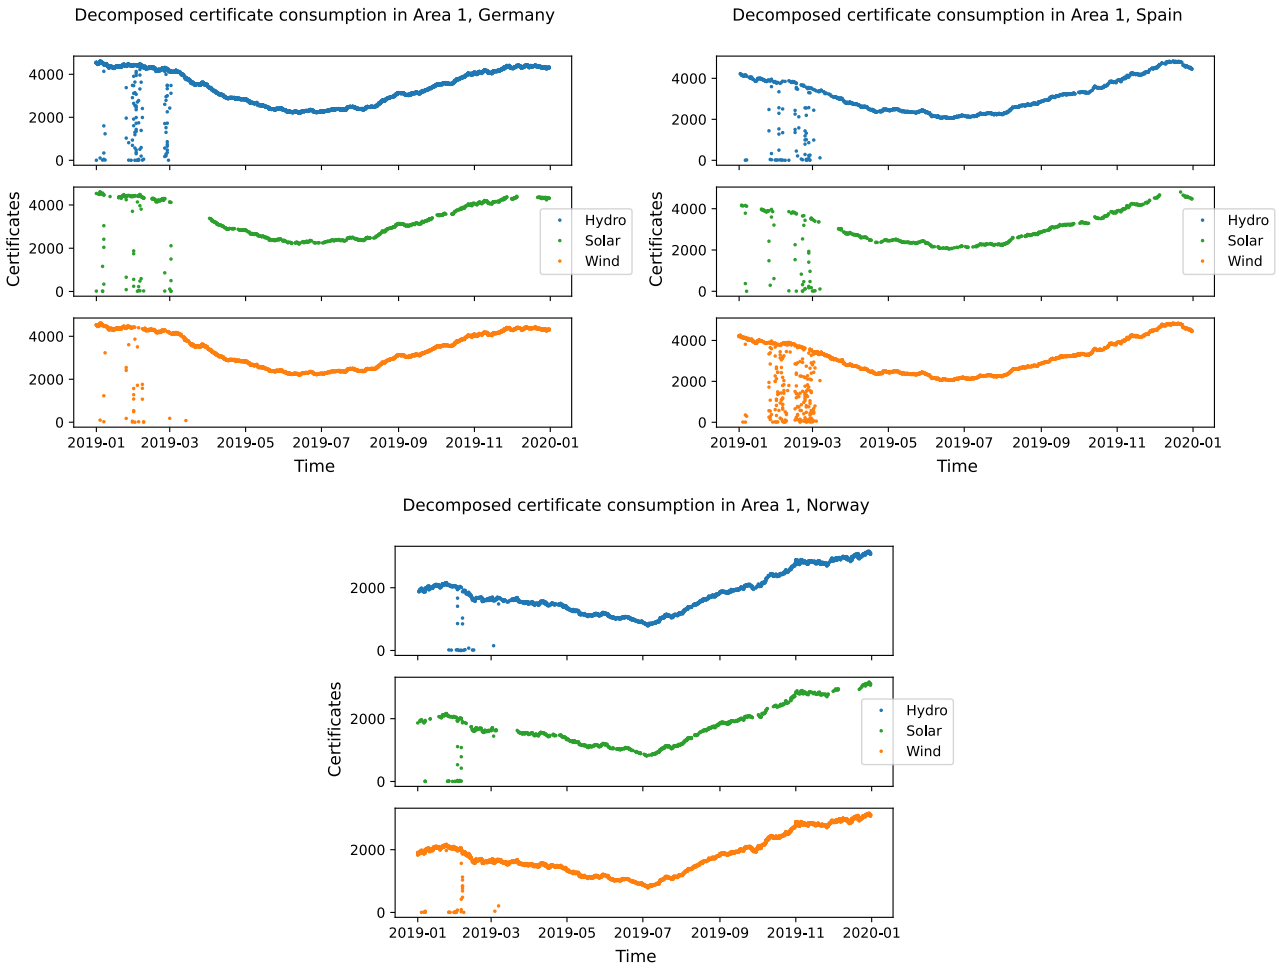

Figure 24: A decomposed view of the certificate consumption plots, showing each certificate group by themselves.

In Figure 25, a 24 hour rolling average of the certificate consumption data is presented. Similarly in Figure 26, a 30 day rolling average is presented. These two graphs better presents the trends of the market. Compared to earlier graphs, it is possible to determine which certificate types dominate each node, and they seem to be either hydro or wind certificates in these example cases. Interestingly, it could appear that the majority certificate type accounts for the seasonal variations in certificate consumption, as especially in Figure 26 the majority certificates follow the seasonal energy consumption pattern. The underlying minority certificate type, which tends to be solar, has a more steady nature, and while it varies it seems to vary around a semi-constant level.

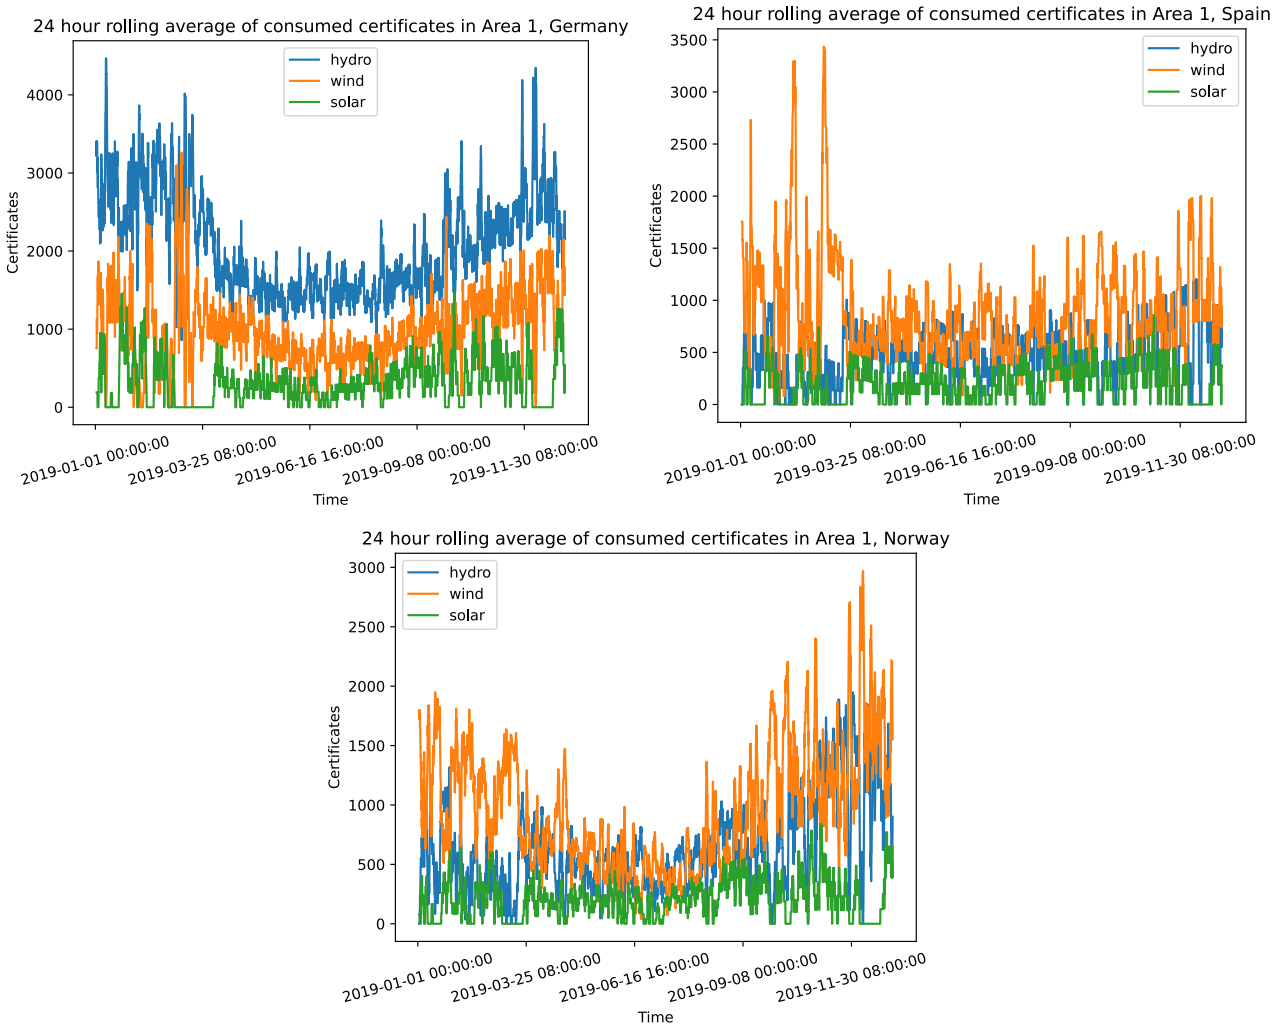

Figure 25: 24 hour rolling average certificate consumption.

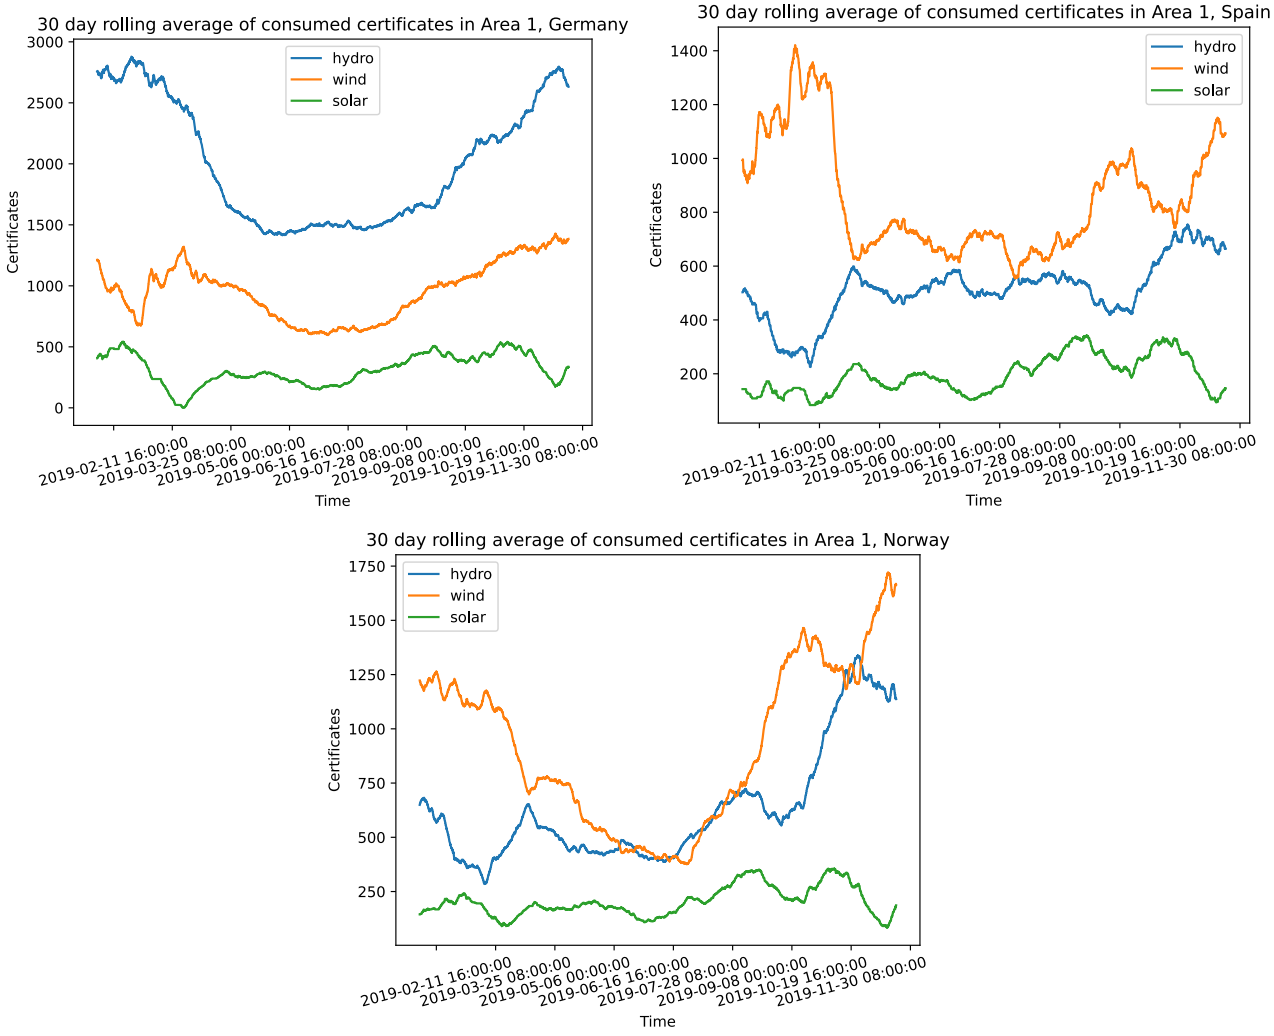

Figure 26: 30 day rolling average certificate consumption.

In Figure 27, the momentarily owned certificate stockpile is shown. Compared to the consumed certificates, each area have distinct certificate types that tend to be stockpiled, which correspond to the certificate types they produce locally. Just as in previous graphs, the certificate deficits found in the first three months can be seen as well. Areas where the owned certificates are zero correspond to timeframes where there are a certificate deficit in the system. Interestingly, the German area seems to reach a constant non-zero plateau of owned solar and hydro certificates. This is a rare result deriving from the network simulations order of operation. Since the nodes place buy orders based on the number of available certificates for sale, and this happens before the trades are executed, the later consumption of the most owned certificate type may vary from which certificate type is bought if the node has a surplus of total certificates. See subsubsection 4.2.4 for the rundown of the order of operations. The other nodes buy certificates to fill their need, and only experience a surplus of the type they produce, when the market is saturated.

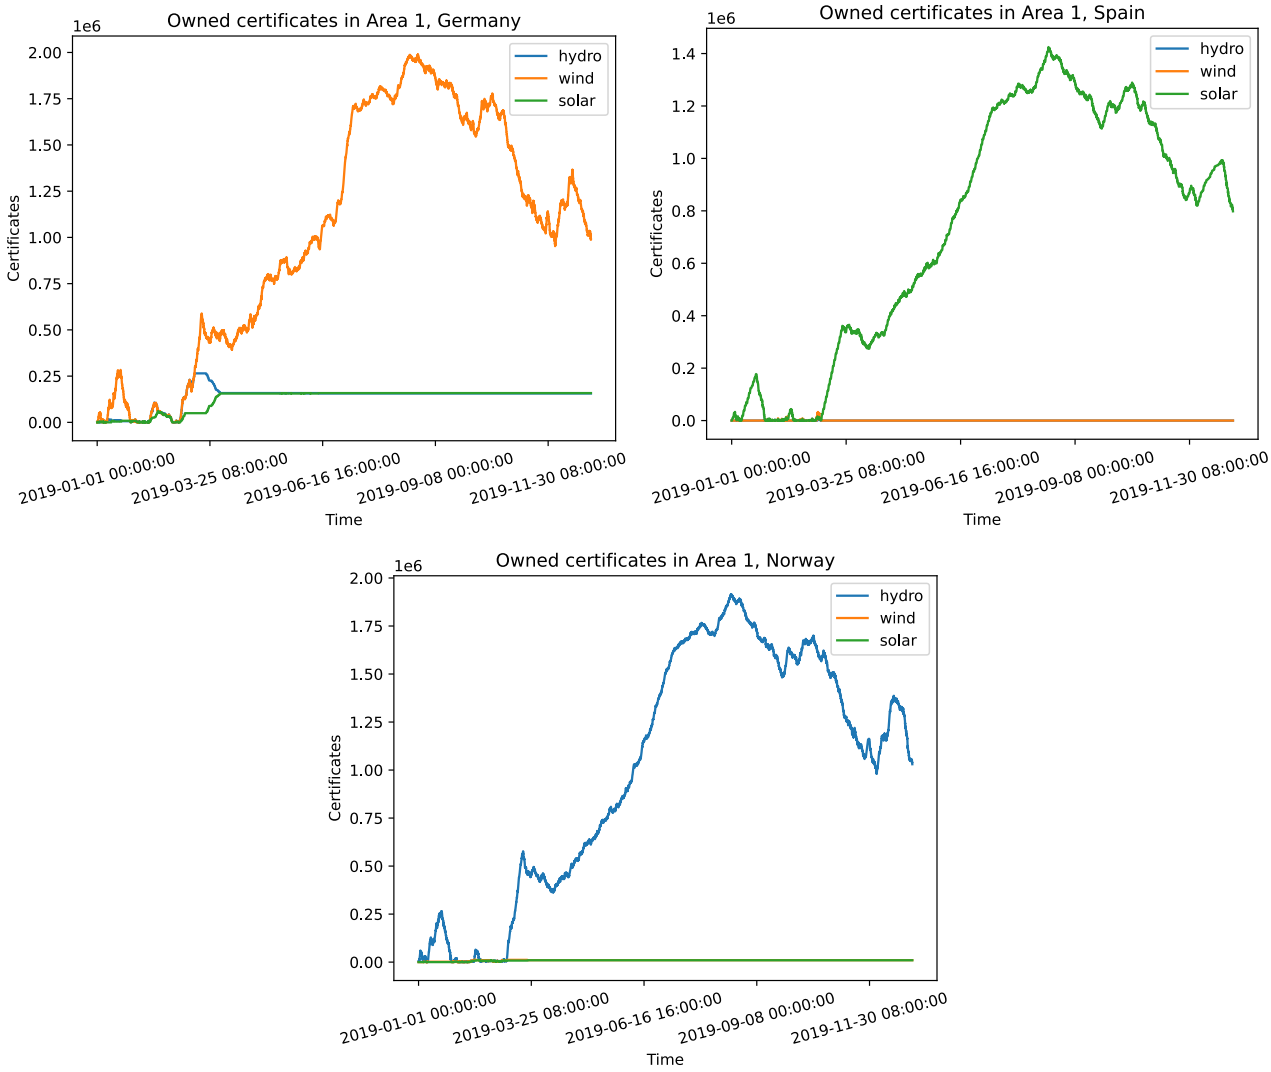

Figure 27: Owned certificates of each area.

Finally, in Figure 28 the wallet balances for each of the fifteen nodes are shown. The balances show the financial development of each node, and as expected some become trade losers and some trade winners.

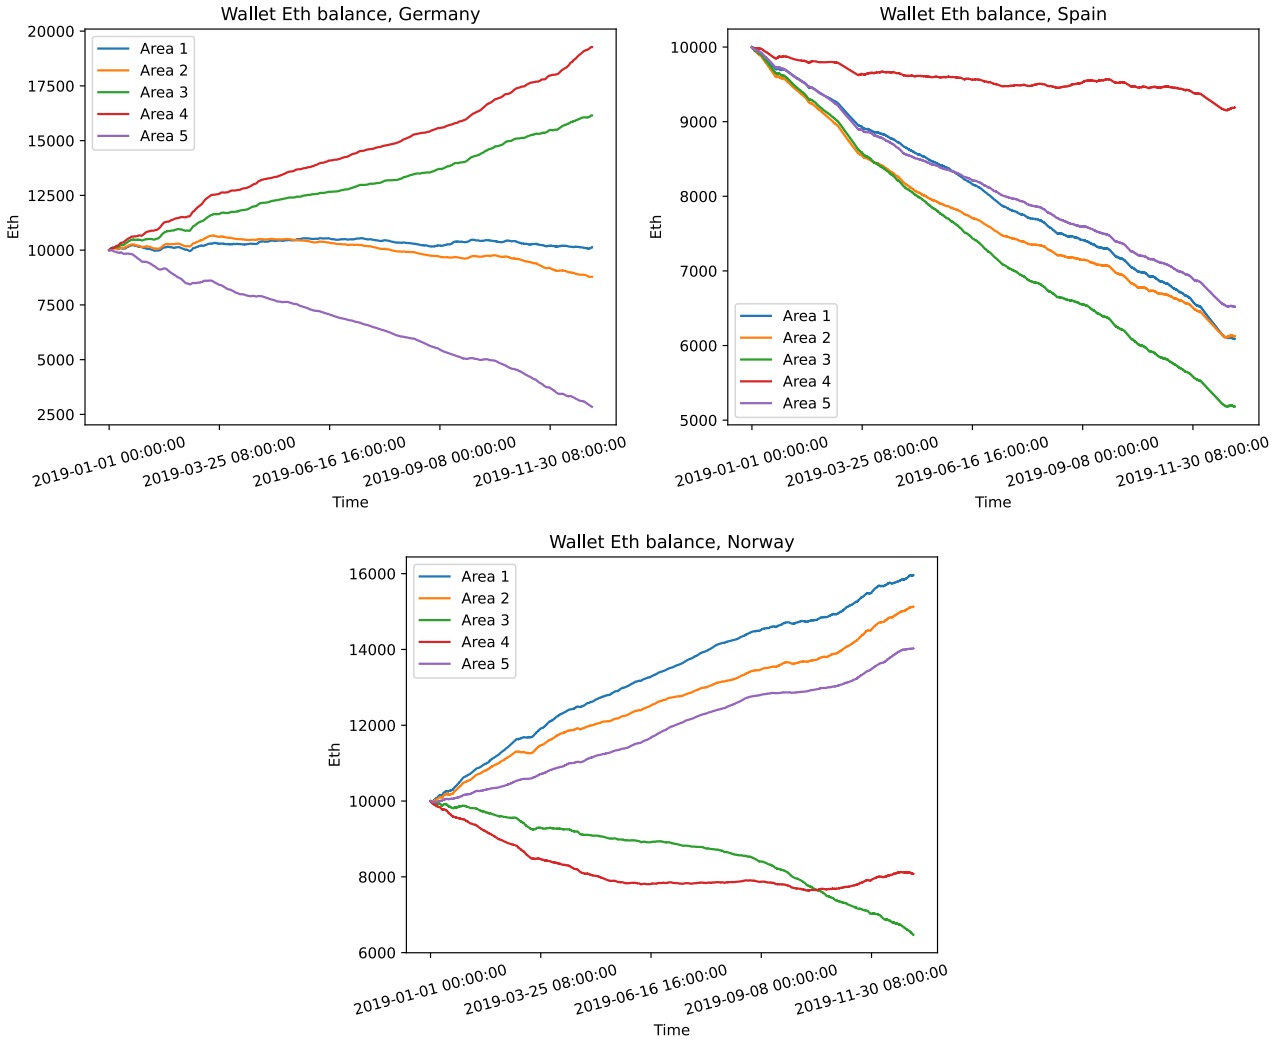

Figure 28: The financial balances of each node.

6 Discussion and Further Work

For this thesis, two on-chain marketplace for energy certificates has been developed and explained. The initial goal was to present a proof of concept, and to evaluate which token standard was more appropriate for large scale certificate trading. The token standards evaluated was the ERC721 standard, which is the common standard for NFTs on Ethereum blockchains, and the ERC1155 standard, which implements batched NFTs/SFTs. The thought was to improve traceability and transparency, while simultaneously provide a single service that could handle near real time trading. After having reviewed the results, there seems to be a trade off between traceability and scalability.

The ERC721 standard provides better traceability, as every certificate can be tracked independently across multiple trades. With unique identifiers for every certificate, two otherwise identical certificates can be bought and sold by the same certificate holder, and still traced back to their origin. Blockchains are intrinsically secured from on-chain double spending errors, but situations where certificates are recorded as spent outside the blockchain, in for example an emission accounting, could lead to certificates being double spent, by both being sold on chain, and consumed off-chain. By using the ERC721 standard, such errors would easily be reversed as the unique token identifier would identify which token has been double spent, and proper reimbursements can be processed. A regulatory body could determine which specific certificate had been double spent, and reverse the transaction.

Comparatively, the ERC1155 standard only provides semi-complete traceability. While it groups the certificates by set characteristics, it is unable to distinguish between two certificates with the same characteristics. This allows for systems where market actors may buy, sell, produce and consume certificates with unique characteristics, but where on/off-chain double spending may lead to problematic situations. Taking the same example as for the ERC721 standard, where a certificate is first consumed off-chain, and then sold on chain, it would not be possible for a regulatory body to determine which exact certificate was the spent one, if the double spender had sold several certificates from the same group.

Considering this difference in transparency, the ERC721 standard appears as the main contender for a well built on-chain certificate marketplace, however it suffers from scalability issues compared to the ERC1155 standard. The time and space complexities presented in subsection 5.3 and subsection 5.4 shows large differences at up-scaled levels. One of the commonly proposed problems with blockchain solutions is the immense and growing storage space occupied by storing information on the blockchain, and forward leaning and sustainable solutions must consider this when presenting scaling solutions. Furthermore, for the sake of this thesis, it was simply not feasible to simulate a large scale node network over a longer period of time, using the ERC721 standard. Hence, the ERC1155 standard proved to be the only sufficient and realistic choice when operating with large quantities of certificates.

The year long simulation did prove successful, and certificates were successfully traded on a blockchain based solution. The solution implemented a simple continuous marketplace, that allowed actors to trade certificates automatically on a decentralized, but still interconnected platform. The simulation assumed issuing intervals of one hour, which allowed for less certificate build up compared to the current standards of either weekly or monthly injections. [49, 47, 48]

While the solution presents a viable proof of concept, there are still areas of the simulation that can be improved. The auction mechanism employed on the marketplace is simple, and only matches viable buy orders and sell orders. It does not optimally match the closest price points as a proper continuous double auction would, instead it matches orders where the ask is lower than the bid in a chronological fashion. While this may provide a good solution in some cases, it does not minimize the differences between bids and asks.

Furthermore the certificate production and consumption data comes from simulated sources. While the data is sufficient to see the marketplace operating, a more thorough solution should simulate a year of certificate trading with real world consumers and producers, and with the corresponding data. This would include more and smaller nodes, which covers the entirety of the EU, and implement semi-random buy-sell trigger times. Currently, the three represented countries buy and sell in a procedural order: Germany, then Spain and then Norway. This creates a certificate preference bias, as the previous country disproportionally floods the market with the corresponding certificate type. A better representation would inject buy and sell orders in varying orders from all nodes.

Currently, the smart contracts are also only deployed on the locally simulated Ganache-cli blockchain. To improve the applicability to the real world, the smart contracts should be deployed on a real private blockchain. This step was under production, and to be done using the Go Ethereum (Geth) tool, however it was postponed in favor of space and speed tests. The Ganache-cli tool provides a life-like substitution, but aspects like network latency and instability are lacking. The lack of proper blockchain deployment means that there might be unforeseen consequences not accounted for when using Ganache-cli.

7 Conclusion

With the rapid energy transition needed to mitigate a climate disaster, the use of energy certificates is becoming more widespread. The European Union has implemented a system that involves Guarantees of Origins, GOs, a specific energy certificate mainly used within the European Economic area. With the digital green shift, GOs are providing much needed energy provenance for consumers, and the trade of them are consequently increasing year over year.

GOs are standarized through the EECS, and at a European scale, the day to day handling are done by the AIB. All nations wishing to issue GOs must have specified domains and corresponding domain protocols. Generally, these domains correspond to each nation, however there are some exceptions. While each domain protocol must adhere to set guidelines by the AIB, they have some freedom of design. For example, the inter-issuing time between GO issuings may vary from domain to domain, with some domains doing weekly issuings, some doing bi-monthly issuings, and some monthly issuings. This means that there is a space for improved standarization and shorter inter-issuing intervals, where each certificate producer has less certificate buildup and greater freedom of trade.

The inter-domain trade of GOs is also shown to be cumbersome, and involve several third party actors. Currently these actors are needed to ensure the legitimacy and correctness of the trades, and to keep statistics and formulate data. However, the need for third parties and separate trading and storing systems implies a space for improved efficiency, if the entire trading process

can be merged into a single tool. This thesis suggests that a DLT ledger may provide the transparency, security and synchronicity needed for such a system.

DLTs encompass a wide range of technologies, but the larges and most prominent is the blockchain. However, there are also a variety of different types and characteristics of blockchains. We showcase the difference between permissioned and permissionless, and private and public blockchains. This distinction is important to consider when assessing which blockchain is the most suitable for a project, as a permissioned and private blockchain will generally be more efficient but less secure, and contrarily a a permissionless and public blockchain will need more resources, but can operate without any inter-party trust.

Blockchains allow for decentralized data storage, good security, transparency and peer to peer interactions. They do, however also have some known downsides. Commonly discussed negatives are the transaction throughput, the energy use related to operation and consensus reaching, and the storage space occupied by the blockchain. As blockchain technologies improve, these drawbacks diminishes, but they should still be considered. The energy consumption issue is mainly related to PoW based blockchains, as for example PoS based blockchains consume only a fraction of the energy of a PoW chain. Similarly for private PoA and PoI blockchains. The transaction throughput can still be a limiting factor, depending on the operations done with the transaction, but newer technologies also show great improvements on the field. With 2.layer technologies, Ethereum estimates a transaction throughput of roughly 5200 transactions per second, and chains like Solana claims a max throughput of up to 65 000 transactions per second, which would be competitive with traditional payment solutions, such as Visa. The issue of storage space is maybe the single properly remaining issue. At the time of writing, the Ethereum blockchain is roughly 650 GB, which is troublesome for small scale nodes, and discourages decentralization of the blockchain.

With all this considered, this thesis has explored the potential for a DLT-based marketplace for energy certificates, and a proof of concept model has been developed and implemented. The goal was to improve the fluidity, trackability, transparency and automaticity of the system. Fundamentally, the idea was to represent each certificate as a Non-Fungible Token on the blockchain, as they provide one to one traceability of each certificate. However, the standard NFT implementation standard (ERC721) was compared to a newer standard for SFTs (ERC1155). Intrinsically, SFTs are less transparent than NFTs, but they provided much better scalability, especially with regards to blockchain space and operational speeds.

A full year simulation was done for fifteen energy producers and consumers, spread across Germany, Spain and Norway. These nodes acted as participants in the developed on-chain marketplace, and conducted trades necessary to fulfill their energy consumption with corresponding energy certificates. The simulations showed a functioning marketplace, and the results showed all available certificates being consumed in times of certificate deficits, a decrease in stored certificates in times where the consumption was greater than the production, and and increase when it was the opposite. In all, it provided a proof of concept showing the potential of a DLT marketplace for certificates. There are, however, components lacking for a fully optimal simulations, and these include among others an improved auction mechanism, better testing data, and an expanded node network.

A Appendix

A.1 Wind simulation graphs.

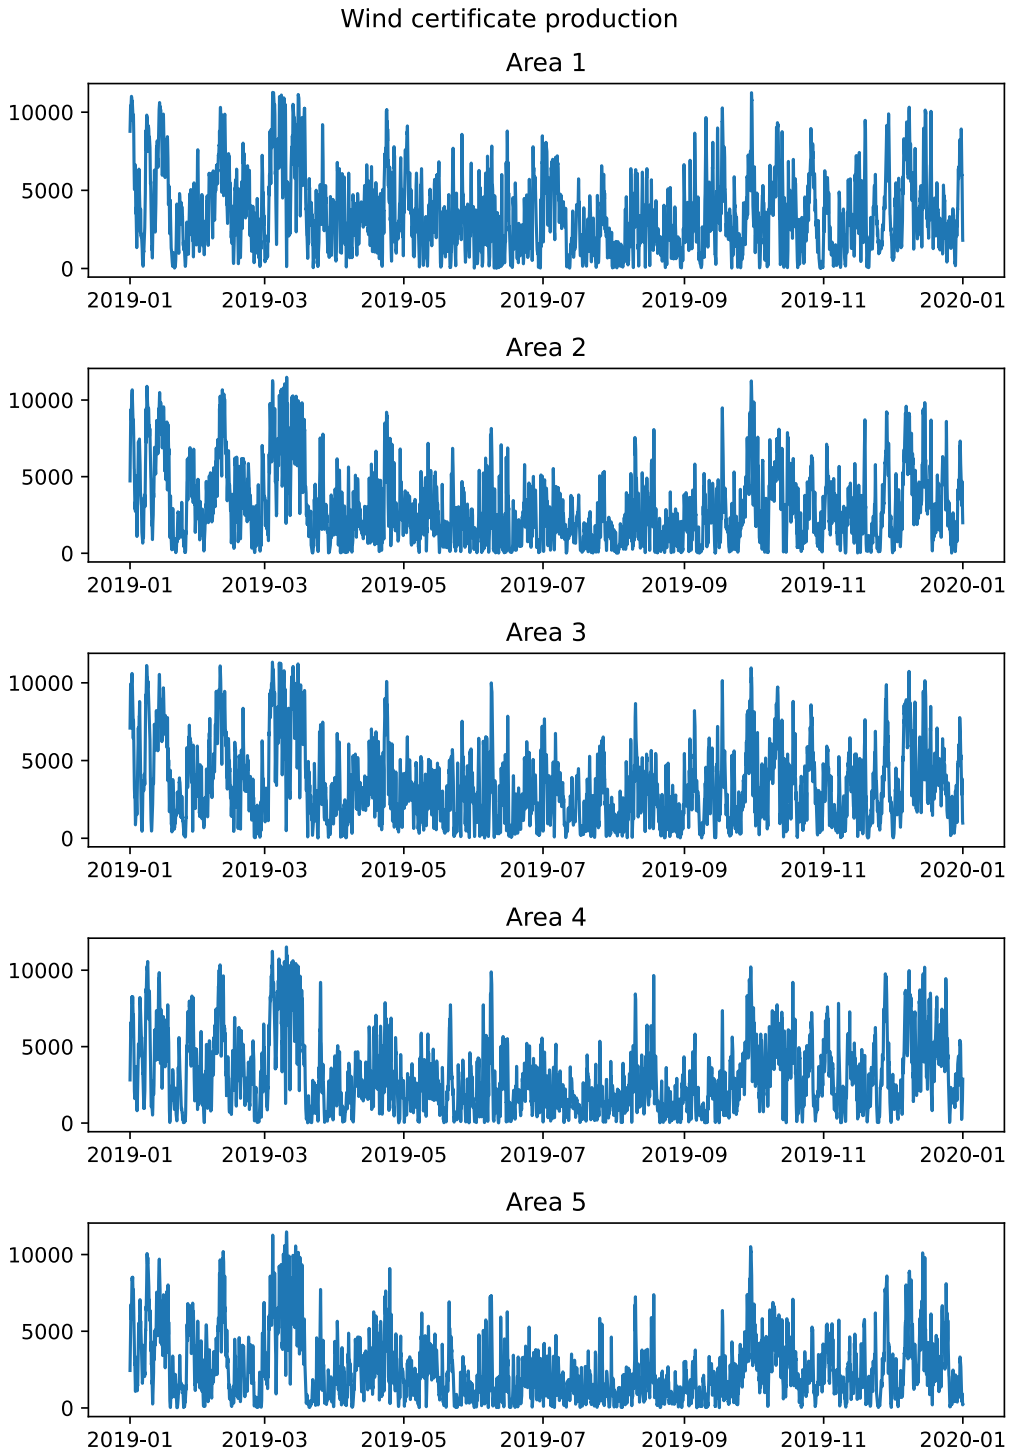

Figure 29: Simulated wind values for all five Areas.

A.2 Solar simulation graphs.

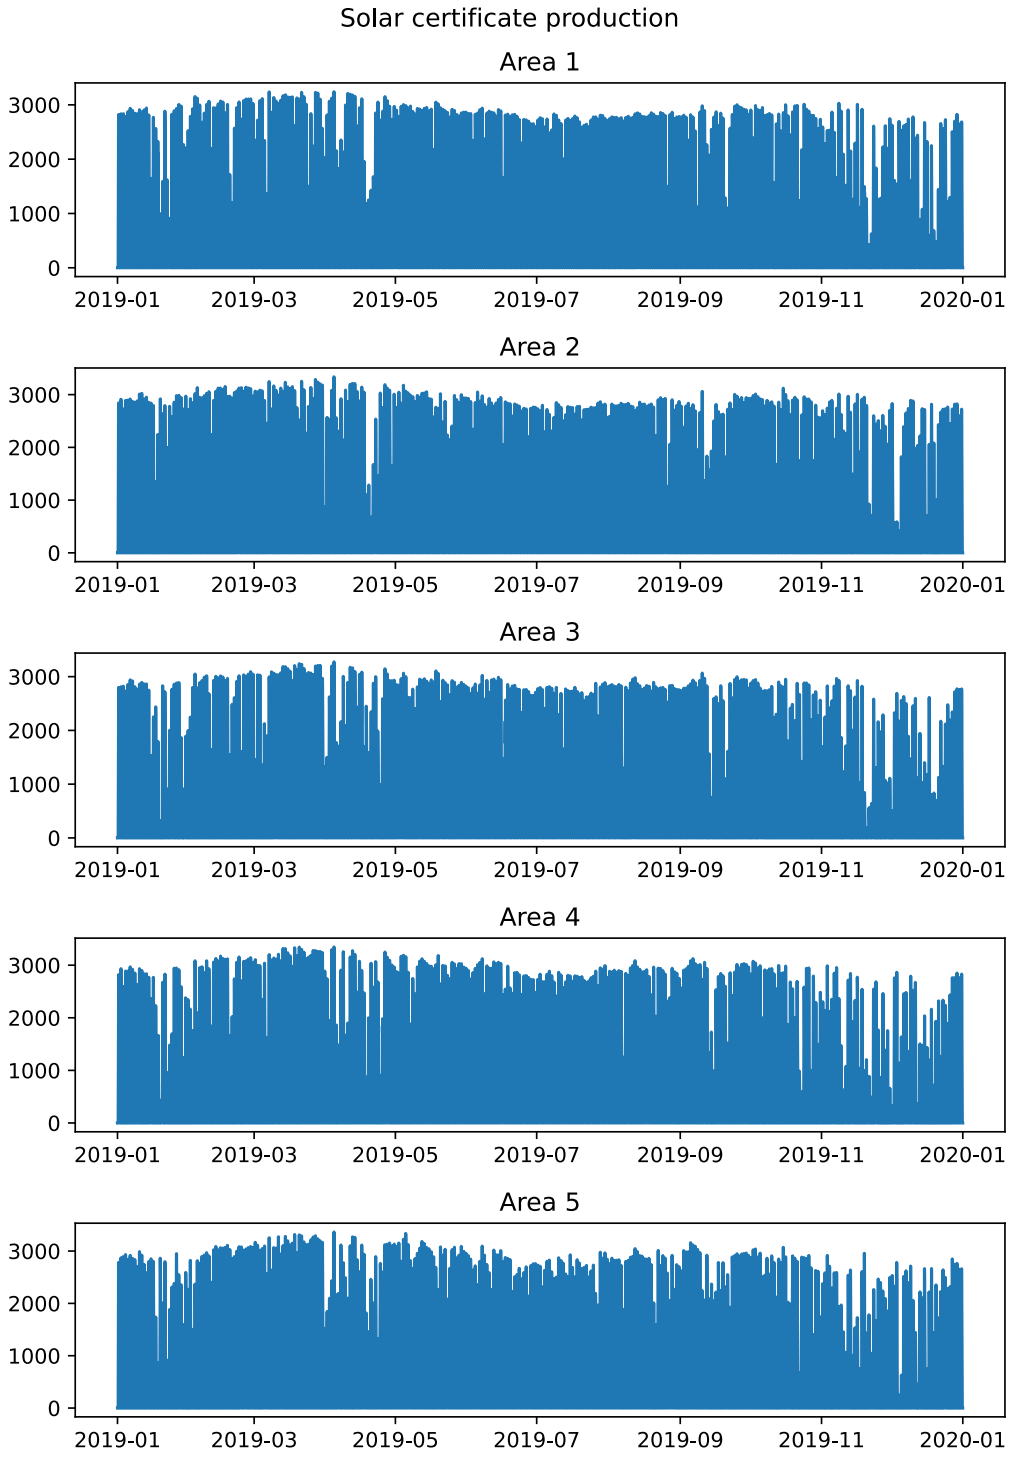

Figure 30: Simulated solar values for all five Areas.

A.3 Hydro simulation graphs.

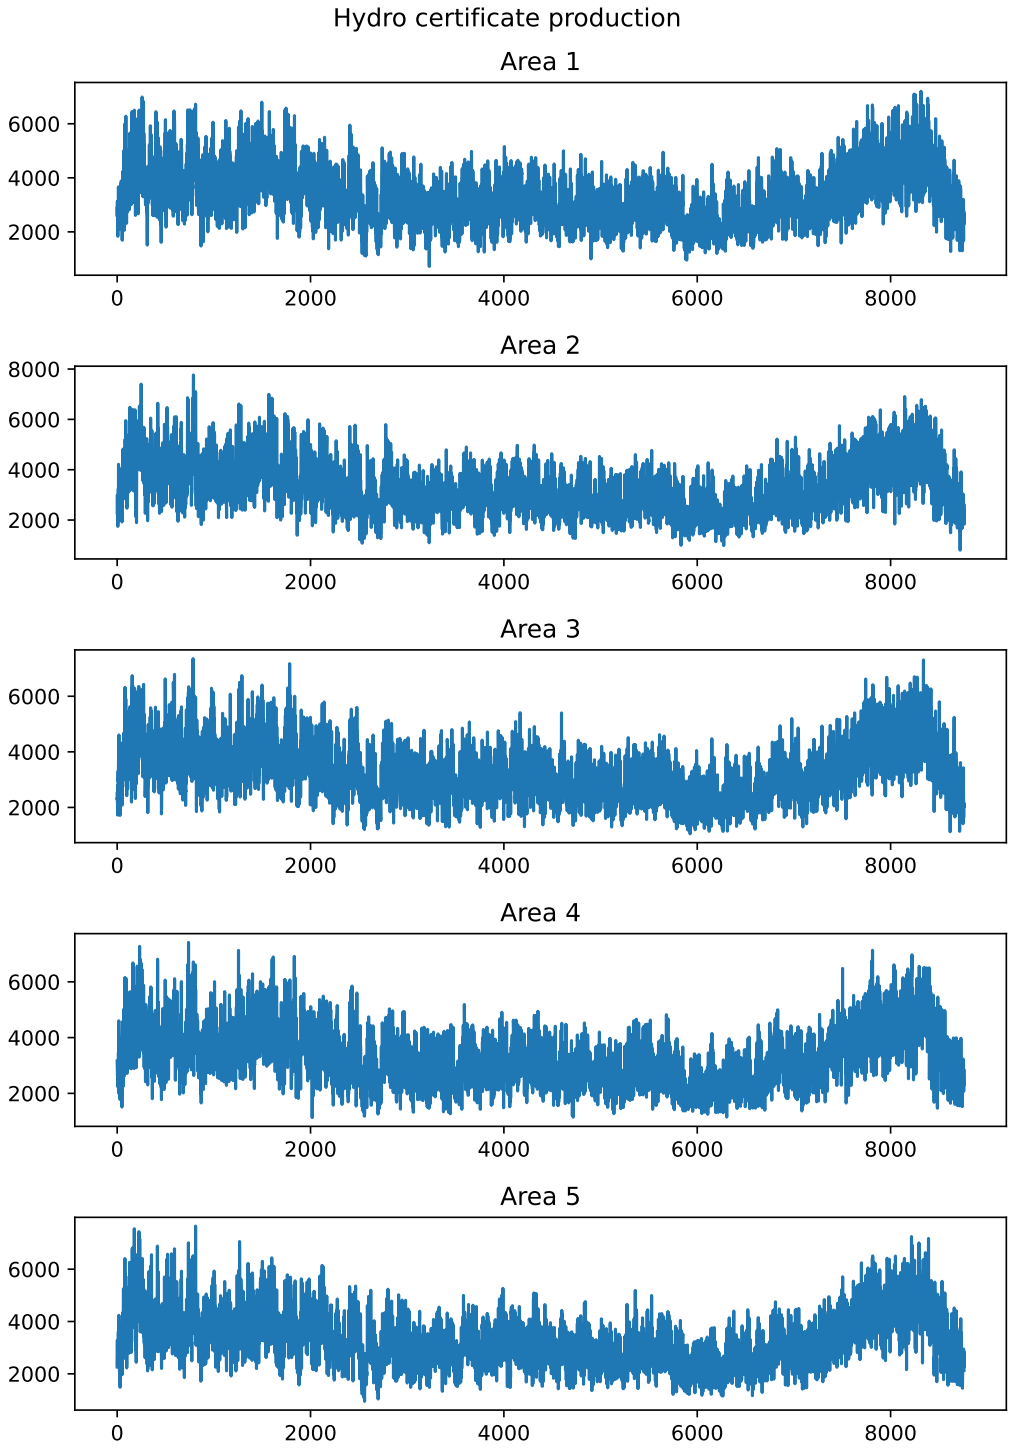

Figure 31: Simulated hydro power values for all five Areas.

References

[1] UNFCCC. *The Paris Agreement | UNFCCC*. https://unfccc.int/process-and-meetings/the-paris-agreement/the-paris-agreement. (Accessed on 07/11/2022).

[2] Ms. Ocasio-Cortez et.al. *BILLS-116hres109ih.pdf*. https://www.congress.gov/116/bills/hres109/BILLS-116hres109ih.pdf. (Accessed on 06/14/2023). Feb. 2019.

[3] Ethereum Foundation. *Ethereum roadmap | ethereum.org*. https://ethereum.org/en/roadmap/. (Accessed on 06/14/2023).

[4] Kristian Astad Dupont. *A DLT-based Architecture for the International Trade of Guarantees of Origin*. (Accessed on 02/13/2023). Dec. 2022.

[5] Hannah Ritchie, Max Roser, and Pablo Rosado. "CO and Greenhouse Gas Emissions". In: *Our World in Data* (2020). https://ourworldindata.org/co2-and-other-greenhouse-gas-emissions.

[6] Fiona Harvey and Damian Carrington. "World is on "highway to climate hell", UN chief warns at Cop27 summit." In: *The Guardian* (Nov. 7, 2022). URL: https://www.theguardian.com/environment/2022/nov/07/cop27-climate-summit-un-secretary-general-antonio-guterres (visited on 11/08/2022).

[7] the United Nations Department of Economic and Social Affairs. *THE 17 GOALS | Sustainable Development*. https://sdgs.un.org/goals. (Accessed on 08/11/2022).

[8] KS Thakur and S Gupta. "Exploration of green shift: shift from trendy marketing to environment friendly green marketing". In: *International Journal of Arts and Commerce* 1.7 (2012), pp. 122–133.

[9] Umit Cali et al. "Introduction to the Digitalization of Power Systems and Markets." In: *Digitalization of Power Markets and Systems Using Energy Informatics*. Cham: Springer International Publishing, 2021, pp. 1–16. ISBN: 978-3-030-83301-5. DOI: 10.1007/978-3-030-83301-5_8. URL: https://doi.org/10.1007/978-3-030-83301-5_8.

[10] Tanveer Ahmad et al. "Artificial intelligence in sustainable energy industry: Status Quo, challenges and opportunities". In: *Journal of Cleaner Production* 289 (2021), p. 125834. ISSN: 0959-6526. DOI: https://doi.org/10.1016/j.jclepro.2021.125834. URL: https://www.sciencedirect.com/science/article/pii/S0959652621000548.

[11] Robert A. Strong. *Jimmy Carter: Domestic Affairs | Miller Center*. https://millercenter.org/president/carter/domestic-affairs. (Accessed on 07/11/2022).

[12] RedClay. *Industry 101 | International Aspects: Deregulation of the Energy Market in the United Kingdom - Red Clay*. https://redclay.com/international-aspects-deregulation-energy-market-united-kingdom/. (Accessed on 07/11/2022). Dec. 2017.

[13] The European Parliament. *Directive 96/92/EC of the European Perliament and of the Council of 19 December 1996 concerning common rules for the internal market in electricity*. https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=CELEX%3A31996L0092. (Accessed on 07/11/2022). Jan. 1997.

[14] Ugur Halden et al. "DLT-based equity crowdfunding on the techno-economic feasibility of solar energy investments". In: *Solar Energy* 227 (2021), pp. 137–150. ISSN: 0038-092X. DOI: https://doi.org/10.1016/j.solener.2021.08.067. URL: https://www.sciencedirect.com/science/article/pii/S0038092X21007313.

[15] Marie Claire Brisbois. "Decentralised energy, decentralised accountability? Lessons on how to govern decentralised electricity transitions from multi-level natural resource gov-

ernance". In: *Global Transitions* 2 (Jan. 2020), pp. 16–25. DOI: 10.1016/j.glt.2020.01.001.

[16] Tobias Boström et al. "The pure PV-EV energy system – A conceptual study of a nationwide energy system based solely on photovoltaics and electric vehicles". In: *Smart Energy* 1 (2021), p. 100001. ISSN: 2666-9552. DOI: https://doi.org/10.1016/j.segy.2021.100001. URL: https://www.sciencedirect.com/science/article/pii/S2666955221000010.

[17] Rajesh Singh et al. "Energy System 4.0: Digitalization of the Energy Sector with Inclination towards Sustainability". In: *Sensors* 22.17 (2022). ISSN: 1424-8220. DOI: 10.3390/s22176619. URL: https://www.mdpi.com/1424-8220/22/17/6619.

[18] Oliver Wagner and Thomas Götz. "Presentation of the 5Ds in Energy Policy: A Policy Paper to Show How Germany Can Regain Its Role as a Pioneer in Energy Policy". In: *Energies* 14.20 (2021). ISSN: 1996-1073. DOI: 10.3390/en14206799. URL: https://www.mdpi.com/1996-1073/14/20/6799.

[19] Memi Motosu and Yasushi Maruyama. "Local acceptance by people with unvoiced opinions living close to a wind farm: A case study from Japan". In: *Energy Policy* 91 (2016), pp. 362–370. ISSN: 0301-4215. DOI: https://doi.org/10.1016/j.enpol.2016.01.018. URL: https://www.sciencedirect.com/science/article/pii/S0301421516300180.

[20] Hannah Ritchie, Max Roser, and Pablo Rosado. "Energy". In: *Our World in Data* (2022). https://ourworldindata.org/energy.

[21] International Renewable Energy Agency. *Renewable Capacity Statistics 2022*. https://www.irena.org/publications/2022/Apr/Renewable-Capacity-Statistics-2022. (Accessed on 12/18/2022). Apr. 2022.

[22] The U.S. Energy Information Administration. *Global Energy Capacity Statistics*. https://www.eia.gov/international/data/world/electricity/electricity-capacity?pd=2&p=0000000000000000000007vo7&u=0&f=A&v=mapbubble&a=-&i=none&vo=value&t=C&g=00000000000000000000000000000000000000000000000001&l=249-ruvvvvvfvtvnvv1vrvvvvfvvvvvvfvvvou20evvvvvvvvvvnvvvs0008&s=315532800000&e=1609459200000&. (Accessed on 12/18/2022). 2022.

[23] Dolf Gielen et al. "The role of renewable energy in the global energy transformation". In: *Energy Strategy Reviews* 24 (2019), pp. 38–50. ISSN: 2211-467X. DOI: https://doi.org/10.1016/j.esr.2019.01.006. URL: https://www.sciencedirect.com/science/article/pii/S2211467X19300082.

[24] The U.S. Energy Information Administration. *Global Energy production statistics*. https://www.eia.gov/international/data/world/electricity/electricity-generation?pd=2&p=0000000000000000000000000000fvu&u=0&f=A&v=mapbubble&a=-&i=none&vo=value&&t=C&g=00000000000000000000000000000000000000000000000001&l=249-ruvvvvvfvtvnvv1vrvvvvfvvvvvvfvvvou20evvvvvvvvvvnvvvs0008&s=315532800000&e=1609459200000. (Accessed on 12/18/2022). 2022.

[25] IRENA. *Renewable Power Generation Costs in 2021*. https://www.irena.org/publications/2022/Jul/Renewable-Power-Generation-Costs-in-2021. (Accessed on 12/18/2022). July 2022.

[26] Christoph Kost. *Study: Levelized Cost of Electricity - Renewable Energy Technologies - Fraunhofer ISE*. https://www.ise.fraunhofer.de/en/publications/studies/cost-of-electricity.html. (Accessed on 12/18/2022). June 2021.

[27] LAZARD. *Lazard.com | Levelized Cost Of Energy, Levelized Cost Of Storage, and Levelized Cost Of Hydrogen*. https://www.lazard.com/perspective/levelized-cost-of-energy-levelized-cost-of-storage-and-levelized-cost-of-hydrogen/. (Accessed on 12/18/2022). Oct. 2021.

[28] Atika Qazi et al. "Towards Sustainable Energy: A Systematic Review of Renewable Energy Sources, Technologies, and Public Opinions". In: *IEEE Access* 7 (2019), pp. 63837–63851. DOI: 10.1109/ACCESS.2019.2906402.

[29] Maria Luisa Di Silvestre et al. "How Decarbonization, Digitalization and Decentralization are changing key power infrastructures". In: *Renewable and Sustainable Energy Reviews* 93 (2018), pp. 483–498. ISSN: 1364-0321. DOI: https://doi.org/10.1016/j.rser.2018.05.068. URL: https://www.sciencedirect.com/science/article/pii/S1364032118304283.

[30] Esther Mengelkamp, Julius Diesing, and Christof Weinhardt. In: *it - Information Technology* 61.2-3 (2019), pp. 101–110. DOI: doi:10.1515/itit-2019-0016. URL: https://doi.org/10.1515/itit-2019-0016.

[31] Esther Mengelkamp et al. "The value of local electricity - A choice experiment among German residential customers". In: *Energy Policy* 130 (2019), pp. 294–303. ISSN: 0301-4215. DOI: https://doi.org/10.1016/j.enpol.2019.04.008. URL: https://www.sciencedirect.com/science/article/pii/S0301421519302447.

[32] Espen Flo Bødal et al. "Challenges in distribution grid with high penetration of renewables". In: *INVADE* 1.0 (2017).

[33] Umit Cali et al. "Cybersecure and scalable, token-based renewable energy certificate framework using blockchain-enabled trading platform". In: *Electrical Engineering* (Nov. 2022). ISSN: 1432-0487. DOI: 10.1007/s00202-022-01688-0. URL: https://doi.org/10.1007/s00202-022-01688-0.

[34] CENELEC. *Smart Grid Reference Architecture*. https://www.cencenelec.eu/media/CEN-CENELEC/AreasOfWork/CEN-CENELEC_Topics/Smart%20Grids%20and%20Meters/Smart%20Grids/reference_architecture_smartgrids.pdf. (Accessed on 12/13/2022). Nov. 2012.

[35] E Holt, J Sumner, and L Bird. "Role of Renewable Energy Certificates in Developing New Renewable Energy Projects". In: (June 2011). DOI: 10.2172/1018490. URL: https://www.osti.gov/biblio/1018490.

[36] Precedence Research. *Renewable Energy Certificate Market Size to Hit USD 100.96 Bn.* https://www.precedenceresearch.com/renewable-energy-certificate-market. (Accessed on 12/10/2022). 2021.

[37] The Rocky Mountain Institute. *Clean Energy 101: The REC Market - RMI*. https://rmi.org/clean-energy-101-the-rec-market/. (Accessed on 12/12/2022). June 2022.

[38] Jane Courtnell. *The Seven Sins of Greenwashing - Green Business Bureau*. https://greenbusinessbureau.com/green-practices/the-seven-sins-of-greenwashing/. (Accessed on 12/12/2022). Dec. 2021.

[39] Sebastião Vieira Freitas Netto et al. "Concepts and forms of greenwashing: a systematic review". In: *Environmental Sciences Europe* 32.1 (Feb. 2020), p. 19. ISSN: 2190-4715. DOI: 10.1186/s12302-020-0300-3. URL: https://doi.org/10.1186/s12302-020-0300-3.

[40] Eric Roston and Ben Elgin. *Companies' Climate Goals in Jeopardy From Flawed Energy Credits*. https://www.bloomberg.com/news/articles/2022-06-09/flawed-

renewable-energy-credits-are-derailing-climate-efforts?leadSource=uverify. (Accessed on 12/12/2022). June 2022.
[41] Lisa Martine Jenkins. *Bitcoin renewable energy monitoring gets a new resource - Protocol*. https://www.protocol.com/bulletins/crypto-mining-renewable-energy-credits. (Accessed on 12/12/2022). May 2022.
[42] Michael Gillenwater. "Redefining RECs—Part 2: Untangling certificates and emission markets". In: *Energy Policy* 36.6 (2008), pp. 2120–2129. ISSN: 0301-4215. DOI: https://doi.org/10.1016/j.enpol.2008.02.019. URL: https://www.sciencedirect.com/science/article/pii/S0301421508000773.
[43] United States Environmental Protection Agency. *State Solar Renewable Energy Certificate Markets | US EPA*. https://www.epa.gov/greenpower/state-solar-renewable-energy-certificate-markets. (Accessed on 11/15/2022). Aug. 2022.
[44] ECOHZ. *Impressive growth in the I-REC market – global Renewable Energy Certificates*. https://www.ecohz.com/news/impressive-growth-in-the-i-rec-market-global-renewable-energy-certificates. (Accessed on 12/12/2022). June 2020.
[45] International REC Standard Fundation. *About Us | I-REC Standard*. https://www.irecstandard.org/about-us/. (Accessed on 12/12/2022). Mar. 2022.
[46] *Domain Protocols | AIB*. https://www.aib-net.org/facts/aib-member-countries-regions/domain-protocols. (Accessed on 11/28/2022). 2022.
[47] Umweltbundesamt (UBA). *EECS Electricity Domain Protocol for Germany*. https://www.aib-net.org/sites/default/files/assets/facts/domain-protocols/Germany%20-%2028-06-2018.pdf. (Accessed on 02/13/2023). 2018.
[48] CNMC. *EECS Electricity Domain Protocol for Spain*. https://www.aib-net.org/sites/default/files/assets/facts/domain-protocols/Spain%20%20-%2007-05-2017.pdf. (Accessed on 04/13/2023). 2018.
[49] Statnett SF. *EECS Electricity Domain Protocol for Norway*. https://www.aib-net.org/sites/default/files/assets/facts/domain-protocols/20180817%20-%20Norway%20-%20Domain%20Protocol.pdf. (Accessed on 02/13/2023). 2018.
[50] William Dixon. *Certificate Auction Platforms - EnergyOrigins.net*. https://energyorigins.net/knowledge-center/auctions/. (Accessed on 11/14/2022). 2022.
[51] Commerg. *Commerg | Greenify your business*. https://commerg.com/. (Accessed on 11/14/2022).
[52] *Global Renewable Energy Trading Network. EECS | I-REC | Elcerts*. https://www.greenpowerhub.com/. (Accessed on 11/14/2022).
[53] EpexSpot. *First pan-European GOs spot auction to take place in September 2022 | EPEX SPOT*. https://www.epexspot.com/en/news/first-pan-european-gos-spot-auction-take-place-september-2022. (Accessed on 11/14/2022). May 2022.
[54] EpexSpot. *Second pan-European spot auction for GOs sees increased market interest | EPEX SPOT*. https://www.epexspot.com/en/news/second-pan-european-spot-auction-gos-sees-increased-market-interest. (Accessed on 11/14/2022). Oct. 2022.
[55] ofGem. *Guarantees of Origin (GoOs) | Ofgem*. https://www.ofgem.gov.uk/environmental-and-social-schemes/renewables-energy-guarantees-origin-rego/energy-suppliers/guarantees-origin-goos. (Accessed on 11/14/2022).
[56] AIB. *AIB monthly statistics*. https://www.aib-net.org/sites/default/files/assets/facts/market%20information/statistics/activity%20statistics/202209%

20AIB%20Statistics%20new%20format%20v1.xlsx. (Accessed on 14/11/2022. Nov. 8, 2022.

[57] Muhammad F. Zia et al. "Microgrid Transactive Energy: Review, Architectures, Distributed Ledger Technologies, and Market Analysis". In: *IEEE Access* 8 (2020), pp. 19410–19432. DOI: 10.1109/ACCESS.2020.2968402.

[58] Umit Cali et al. "Foundations of Distributed Ledger Technology". In: *Digitalization of Power Markets and Systems Using Energy Informatics*. Cham: Springer International Publishing, 2021, pp. 169–195. ISBN: 978-3-030-83301-5. DOI: 10.1007/978-3-030-83301-5_8. URL: https://doi.org/10.1007/978-3-030-83301-5_8.

[59] Umit Cali et al. "Energy Systems Meet with Blockchain Technology". In: *Digitalization of Power Markets and Systems Using Energy Informatics*. Cham: Springer International Publishing, 2021, pp. 197–216. ISBN: 978-3-030-83301-5. DOI: 10.1007/978-3-030-83301-5_9. URL: https://doi.org/10.1007/978-3-030-83301-5_9.

[60] Ammar Odeh, Ismail Keshta, and Qasem Abu Al-Haija. "Analysis of Blockchain in the Healthcare Sector: Application and Issues". In: *Symmetry* 14.9 (2022), p. 1760.

[61] Seyed Mojtaba Hosseini Bamakan et al. "A Decentralized Framework for Patents and Intellectual Property as NFT in Blockchain Networks". In: (2021).

[62] Surya Viswanathen and Aakash Shah. *The Scalability Trilemma in Blockchain | by NeonVest | Medium*. https://aakash-111.medium.com/the-scalability-trilemma-in-blockchain-75fb57f646df. (Accessed on 13/10/2022). Oct. 2018.

[63] Arvind Narayanan et al. *Bitcoin and Cryptocurrency Technologies: A Comprehensive Introduction*. 41 William Street, Princeton, New Jersey 08540, United Kingdom: Princeton University Press, 2016. ISBN: 978-0-691-17169-2.

[64] Junjie Hu and Ke Liu. "Raft consensus mechanism and the applications". In: *Journal of Physics: Conference Series* 1544.1 (May 2020), p. 012079. DOI: 10.1088/1742-6596/1544/1/012079. URL: https://dx.doi.org/10.1088/1742-6596/1544/1/012079.

[65] Nikos Chondros, Konstantinos Kokordelis, and Mema Roussopoulos. "On the Practicality of Practical Byzantine Fault Tolerance". In: *Middleware 2012*. Ed. by Priya Narasimhan and Peter Triantafillou. Berlin, Heidelberg: Springer Berlin Heidelberg, 2012, pp. 436–455. ISBN: 978-3-642-35170-9.

[66] Christian Gorenflo et al. "FastFabric: Scaling Hyperledger Fabric to 20,000 Transactions per Second". In: May 2019, pp. 455–463. DOI: 10.1109/BLOC.2019.8751452.

[67] John R. Douceur. "The Sybil Attack". In: *Peer-to-Peer Systems*. Ed. by Peter Druschel, Frans Kaashoek, and Antony Rowstron. Berlin, Heidelberg: Springer Berlin Heidelberg, 2002, pp. 251–260. ISBN: 978-3-540-45748-0.

[68] Gianmaria Del Monte, Diego Pennino, and Maurizio Pizzonia. "Scaling Blockchains without Giving up Decentralization and Security: A Solution to the Blockchain Scalability Trilemma". In: *Proceedings of the 3rd Workshop on Cryptocurrencies and Blockchains for Distributed Systems*. CryBlock '20. London, United Kingdom: Association for Computing Machinery, 2020, pp. 71–76. ISBN: 9781450380799. DOI: 10.1145/3410699.3413800. URL: https://doi.org/10.1145/3410699.3413800.

[69] Vitalik Buterin. *Why sharding is great: demystifying the technical properties*. https://vitalik.ca/general/2021/04/07/sharding.html. (Accessed on 11/10/2022). Apr. 2021.

[70] Luis Oliveira et al. “To Token or not to Token: Tools for Understanding Blockchain Tokens”. In: Oct. 2018. URL: https://www.researchgate.net/publication/328162731_To_Token_or_not_to_Token_Tools_for_Understanding_Blockchain_Tokens.

[71] Lorien Gabel. *(Mis)Understanding Yield and Inflation in Proof of Stake Networks*. https://www.figment.io/resources/misunderstanding-yield-and-inflation-in-proof-of-stake-networks. (Accessed on 12/10/2022). May 2019.

[72] Energy Labs. *EnergoLabs whitepaperV1.1.3*. https://cryptorating.eu/whitepapers/Energo/[EnergoLabs]%20whitepaperV1.1.3.pdf. (Accessed on 12/10/2022). 2017.

[73] *Exergy Token: The first energy marketplace for the new energy consumer*. http://www.truevaluemetrics.org/DBpdfs/Initiatives/Exergy/Exergy-2018-Executive-Summary.pdf. (Accessed on 12/10/2022). 2018.

[74] Usman W. Choman. *Non-Fungible Tokens: Blockchains, Scarcity, and Value*. (Accessed on 12/10/2022). Mar. 2021. URL: http://dx.doi.org/10.2139/ssrn.3822743.

[75] Hong Bao and David Roubaud. “Non-Fungible Token: A Systematic Review and Research Agenda”. In: *Journal of Risk and Financial Management* 15.5 (2022). ISSN: 1911-8074. DOI: 10.3390/jrfm15050215. URL: https://www.mdpi.com/1911-8074/15/5/215.

[76] Qin Wang et al. “Non-fungible token (NFT): Overview, evaluation, opportunities and challenges”. In: *arXiv preprint arXiv:2105.07447* (2021).

[77] Paul Wackerow. *ERC-721 Non-Fungible Token Standard | ethereum.org*. https://ethereum.org/en/developers/docs/standards/tokens/erc-721/. (Accessed on 12/10/2022). Aug. 2022.

[78] Witek Radomski et al. *ERC-1155: Multi Token Standard*. https://eips.ethereum.org/EIPS/eip-1155. (Accessed on 05/09/2023). 2018.

[79] Gabriel Antonio F. Rebello et al. “A security and performance analysis of proof-based consensus protocols”. In: *Annals of Telecommunications* 77.7 (Aug. 2022), pp. 517–537. ISSN: 1958-9395. DOI: 10.1007/s12243-021-00896-2. URL: https://doi.org/10.1007/s12243-021-00896-2.

[80] Truong Nguyen and Kyungbaek Kim. “A survey about consensus algorithms used in Blockchain”. In: *Journal of Information Processing Systems* 14 (Jan. 2018), pp. 101–128. DOI: 10.3745/JIPS.01.0024.

[81] Benedict George. *What Is a Consensus Mechanism?* https://www.coindesk.com/learn/what-is-a-consensus-mechanism/. (Accessed on 15/10/2022). May 2022.

[82] Sarvagya Bansai. *The Story of those that DIDN’T make it — Stales, Uncles and . . . Orphans? | by Sarvagya Bansal | tech@iiit-gwalior | Medium*. https://medium.com/tech-iiitg/the-story-of-those-that-didnt-make-it-stales-uncles-and-orphans-cea80de7c549. (Accessed on 12/19/2022). May 2020.

[83] CoinBundle Team. *Soft Forks and Hard Forks. Forking explained | by CoinBundle Team | CoinBundle | Medium*. https://medium.com/coinbundle/soft-forks-and-hard-forks-d646e0a21bb6. (Accessed on 12/19/2022). Sept. 2018.

[84] Cong T. Nguyen et al. “Proof-of-Stake Consensus Mechanisms for Future Blockchain Networks: Fundamentals, Applications and Opportunities”. In: *IEEE Access* 7 (2019), pp. 85727–85745. DOI: 10.1109/ACCESS.2019.2925010.

[85] Bhaskar Kashyap. *Proof-of-stake (PoS) | ethereum.org*. https://ethereum.org/en/developers/docs/consensus-mechanisms/pos/#pos-and-security. (Accessed on 06/13/2023). May 2023.

[86] Yossi Gilad et al. “Algorand: Scaling Byzantine Agreements for Cryptocurrencies”. In: Oct. 2017, pp. 51–68. DOI: 10.1145/3132747.3132757.

[87] Umededoteth. *Optimistic Rollups | ethereum.org*. https://ethereum.org/en/developers/docs/scaling/optimistic-rollups/#scaling-ethereum-with-optimistic-rollups. (Accessed on 06/15/2023). May 2023.

[88] *Solana | Web3 Infrastructure for Everyone*. https://solana.com/. (Accessed on 12/19/2022).

[89] *Solana Compass: Solana Staking, TPS Performance Statistics, Validator + More*. https://solanacompass.com/#:~:text=High%20throughput%3A%20Solana%20can%20handle,the%20fastest%20blockchains%20in%20existence.. (Accessed on 06/15/2023).

[90] Miguel Castro, Barbara Liskov, et al. “Practical byzantine fault tolerance”. In: *OsDI*. Vol. 99. 1999. 1999, pp. 173–186.

[91] Rüdiger Reischuk. “A new solution for the Byzantine generals problem”. In: *Information and Control* 64.1 (1985). International Conference on Foundations of Computation Theory, pp. 23–42. ISSN: 0019-9958. DOI: https://doi.org/10.1016/S0019-9958(85)80042-5. URL: https://www.sciencedirect.com/science/article/pii/S0019995885800425.

[92] Hyperledger Foundation. *Introduction — hyperledger-fabricdocs main documentation*. https://hyperledger-fabric.readthedocs.io/en/release-2.2/whatis.html#pluggable-consensus. (Accessed on 09/06/2022). 2022.

[93] Vojislav B. Mišić, Jelena Mišić, and Xiaolin Chang. “Coping with smartly malicious leaders: PBFT with arbitration for blockchain-based IoT applications”. In: *2021 IEEE Global Communications Conference (GLOBECOM)*. 2021, pp. 1–6. DOI: 10.1109/GLOBECOM46510.2021.9685655.

[94] *Bitcoin Gold (BTG) was 51% attacked*. https://gist.github.com/metalicjames/71321570a105940529e709651d0a9765. (Accessed on 31/10/2022).

[95] *exp-attack.md*. https://gist.github.com/metalicjames/\protect\@normalcr\relax01222049f95f85df8c0eb253de54848b. (Accessed on 31/10/2022).

[96] *lcc-attacks.md*. https://gist.github.com/metalicjames/\protect\@normalcr\relax82a49f8afa87334f929881e55ad4ffd7. (Accessed on 31/10/2022).

[97] *Pool Stats - BTC.com*. https://btc.com/stats/pool. (Accessed on 31/10/2022).

[98] Rachel McIntosh. *2 Bitcoin Cash Mining Pools Organized 51% Attack to Thwart Hacker | Finance Magnates*. https://www.financemagnates.com/cryptocurrency/news/2-bitcoin-cash-mining-pools-organized-51-attack-to-thwart-hacker/. (Accessed on 31/10/2022). May 2019.

[99] Evangelos Deirmentzoglou, Georgios Papakyriakopoulos, and Constantinos Patsakis. “A Survey on Long-Range Attacks for Proof of Stake Protocols”. In: *IEEE Access* 7 (2019), pp. 28712–28725. DOI: 10.1109/ACCESS.2019.2901858.

[100] aul Wackerow. *Weak subjectivity | ethereum.org*. https://ethereum.org/en/developers/docs/consensus-mechanisms/pos/weak-subjectivity/. (Accessed on 31/10/2022). Aug. 2022.

[101] Sarah Azouvi and Marko Vukolić. *Pikachu: Securing PoS Blockchains from Long-Range Attacks by Checkpointing into Bitcoin PoW using Taproot*. 2022. DOI: 10.48550/ARXIV.2208.05408. URL: https://arxiv.org/abs/2208.05408.

[102] CoinMarketCap. *Cryptocurrency Prices, Charts And Market Capitalizations | CoinMarketCap*. https://coinmarketcap.com/. (Accessed on 10/10/2022).

[103] *History and Forks of Ethereum | ethereum.org*. https://ethereum.org/en/history/#homestead. (Accessed on 10/10/2022).
[104] *Introduction to Smart Contracts — Solidity 0.8.17 documentation*. https://docs.soliditylang.org/en/v0.8.17/introduction-to-smart-contracts.html?highlight=storage#storage-memory-and-the-stack. (Accessed on 10/10/2022).
[105] Minimalsm. *Ethereum Virtual Machine (EVM) | ethereum.org*. https://ethereum.org/en/developers/docs/evm/. (Accessed on 05/19/2023). Jan. 2023.
[106] *Ethereum Whitepaper | ethereum.org*. https://ethereum.org/en/whitepaper/. (Accessed on 10/10/2022).
[107] Chris Dannen. In: *Introducing Ethereum and Solidity | Foundations of Cryptocurrency and Blockchain Programming for Beginners*. Springer International Publishing, 2017, pp. 69–88. ISBN: 978-1-4842-2534-9. DOI: 10.1007/978-1-4842-2535-6. URL: https://link.springer.com/content/pdf/10.1007/978-1-4842-2535-6.pdf.
[108] Chris Dannen. In: *Introducing Ethereum and Solidity | Foundations of Cryptocurrency and Blockchain Programming for Beginners*. Springer International Publishing, 2017, pp. 47–68. ISBN: 978-1-4842-2534-9. DOI: 10.1007/978-1-4842-2535-6. URL: https://link.springer.com/content/pdf/10.1007/978-1-4842-2535-6.pdf.
[109] Dr. Gavin Wood. *Ethereum: A secure and decentralized generalized transaction ledger Berlin Version*. https://ethereum.github.io/yellowpaper/paper.pdf. (Accessed on 04/11/2023). Oct. 2022.
[110] Victor Genin. *Solidity Gas Optimisation*. https://deeprnd.blogspot.com/solidity-gas-optimisation. (Accessed on 12/02/2022). Jan. 2022.
[111] Abhishek Vispute and Thomas Loimayr. *Solidity Gas Optimizations Cheat Sheet*. https://0xmacro.com/blog/solidity-gas-optimizations-cheat-sheet/. (Accessed on 12/19/2022). July 2022.
[112] *Ethereum vision | ethereum.org*. https://ethereum.org/en/upgrades/vision/. (Accessed on 10/10/2022).
[113] Vinay Gupta. *The Ethereum Launch Process | Ethereum Foundation Blog*. https://blog.ethereum.org/2015/03/03/ethereum-launch-process. (Accessed on 10/10/2022). Mar. 2015.
[114] *The Beacon Chain | ethereum.org*. https://ethereum.org/en/upgrades/beacon-chain/. (Accessed on 10/10/2022). Oct. 2022.
[115] Carl Beekhuizen. *Ethereum's energy usage will soon decrease by ~99.95% | Ethereum Foundation Blog*. https://blog.ethereum.org/2021/05/18/country-power-no-more. (Accessed on 10/10/2022). May 2021.
[116] Paul Wackerow. *Optimistic Rollups | ethereum.org*. https://ethereum.org/en/developers/docs/scaling/optimistic-rollups/. (Accessed on 11/10/2022). Sept. 2022.
[117] Ethereum Foundation. *Danksharding | ethereum.org*. https://ethereum.org/en/roadmap/danksharding/. (Accessed on 06/14/2023). June 2023.
[118] Miguel Muñoz Ortiz, Lisa Kvalbein, and Lars Hellemo. "Evaluation of open photovoltaic and wind production time series for Norwegian locations". In: *Energy* 236 (2021), p. 121409. ISSN: 0360-5442. DOI: https://doi.org/10.1016/j.energy.2021.121409. URL: https://www.sciencedirect.com/science/article/pii/S0360544221016571.
[119] Stefan Pfenninger and Iain Staffell. "Long-term patterns of European PV output using 30 years of validated hourly reanalysis and satellite data". In: *Energy* 114 (2016), pp. 1251–

1265. ISSN: 0360-5442. DOI: https://doi.org/10.1016/j.energy.2016.08.060. URL: https://www.sciencedirect.com/science/article/pii/S0360544216311744.

[120] Thomas Huld et al. "Mapping the performance of PV modules, effects of module type and data averaging". In: *Solar Energy* 84.2 (2010), pp. 324–338. ISSN: 0038-092X. DOI: https://doi.org/10.1016/j.solener.2009.12.002. URL: https://www.sciencedirect.com/science/article/pii/S0038092X0900293X.

[121] Francisco G. Montoya, Maria J. Aguilera, and Francisco Manzano-Agugliaro. "Renewable energy production in Spain: A review". In: *Renewable and Sustainable Energy Reviews* 33 (2014), pp. 509–531. ISSN: 1364-0321. DOI: https://doi.org/10.1016/j.rser.2014.01.091. URL: https://www.sciencedirect.com/science/article/pii/S1364032114001178.

[122] Iain Staffell and Stefan Pfenninger. "Using bias-corrected reanalysis to simulate current and future wind power output". In: *Energy* 114 (2016), pp. 1224–1239. ISSN: 0360-5442. DOI: https://doi.org/10.1016/j.energy.2016.08.068. URL: https://www.sciencedirect.com/science/article/pii/S0360544216311811.

[123] Iain Staffell and Stefan Pfenninger. *Supplementary Material to Using Bias-Corrected Reanalysis to Simulate Current and Future Wind Power Output*. https://ars.els-cdn.com/content/image/1-s2.0-S0360544216311811-mmc1.pdf. (Accessed on 04/21/2023). 2016.

[124] NASA. *MERRA-2 FAQ*. https://gmao.gsfc.nasa.gov/reanalysis/MERRA-2/FAQ/. (Accessed on 04/21/2023).

[125] Global Modeling and Assimilation Office (GMAO), Goddard Earth Sciences Data and Information Services Center (GES DISC). *MERRA-2 $inst1_2d_asm_Nx$: $2d, 1-Hourly, Instantaneous, Single-Level, Assimilation, Single-LevelDiagnosticsV5.12.4$*. accessed 03.05.2023. 2015. DOI: 10.5067/3Z173KIE2TPD.

[126] Vestas. *Vestas V47 - 660,00 kW - Wind turbine*. https://en.wind-turbine-models.com/turbines/13-vestas-v47. (Accessed on 05/03/2023).

[127] STATTNETT. *Data from the power system | Statnett*. https://www.statnett.no/en/for-stakeholders-in-the-power-industry/data-from-the-power-system/#production-and-consumption. (Accessed on 05/08/2023).

[128] Energifakta Norge. *Electricity production - Energifakta Norge*. https://energifaktanorge.no/en/norsk-energiforsyning/kraftproduksjon/. (Accessed on 04/10/2023). May 2021.

[129] SMARD. *SMARD | The electricity market in 2022*. https://www.smard.de/page/en/topic-article/5892/209668. (Accessed on 05/08/2023). 2022.

[130] Sequence. *Introduction to Semi-Fungible Tokens (SFTs) | by Sequence | Sequence | Medium*. https://medium.com/0xsequence/introduction-to-semi-fungible-tokens-sfts-9275bbee4498. (Accessed on 04/12/2023). Feb. 2023.

[131] the Ethereum Foundation. *gm — web3.py 6.2.0 documentation*. https://web3py.readthedocs.io/en/stable/. (Accessed on 05/14/2023).

[132] *Ganache: A tool for creating a local blockchain for fast Ethereum development*. https://github.com/trufflesuite/ganache. (Accessed on 05/14/2023).

[133] Bernard Peh. *Solidity Bytecode and Opcode Basics | by Bernard Peh | Medium*. https://medium.com/@blockchain101/solidity-bytecode-and-opcode-basics-672e9b1a88c2. (Accessed on 05/14/2023). Sept. 2017.

[134] Piper Merriam. *py-solc-x — py-solc-x 1.1.1 documentation*. https://solcx.readthedocs.io/en/latest/. (Accessed on 05/14/2023).

[135] Manal Patel. *Ethereum Series — Understanding Nonce | by Manan Patel | The Startup | Medium*. https://medium.com/swlh/ethereum-series-understanding-nonce-3858194b39bf. (Accessed on 05/14/2023). Aug. 2019.

[136] Stefa Hinum. *Acer Aspire S13 S5-371-77A7 - Notebookcheck.net External Reviews*. https://www.notebookcheck.net/Acer-Aspire-S13-S5-371-77A7.178807.0.html. (Accessed on 05/18/2023). Sept. 2016.

[137] Bundesverband WindEnergie. *Statistics Germany | BWE e.V.* https://www.wind-energie.de/english/statistics/statistics-germany/. (Accessed on 04/10/2023).

[138] *Solar PV capacity Spain 2010-2022 | Statista*. https://www.statista.com/statistics/1003707/installed-solar-pv-capacity-in-spain/. (Accessed on 04/10/2023).

[139] Kristian Elster and Peter Svaar. *Ekspert: – Ligger an til gratis strøm resten av året – NRK Norge – Oversikt over nyheter fra ulike deler av landet*. https://www.nrk.no/norge/ekspert_-_ligger-an-til-gratis-strom-resten-av-aret-1.16232738. (Accessed on 05/30/2023). Dec. 2022.

[140] Paul Wackerow. *Introduction to smart contracts | ethereum.org*. https://ethereum.org/en/developers/docs/smart-contracts/. (Accessed on 06/05/2023). Sept. 2022.

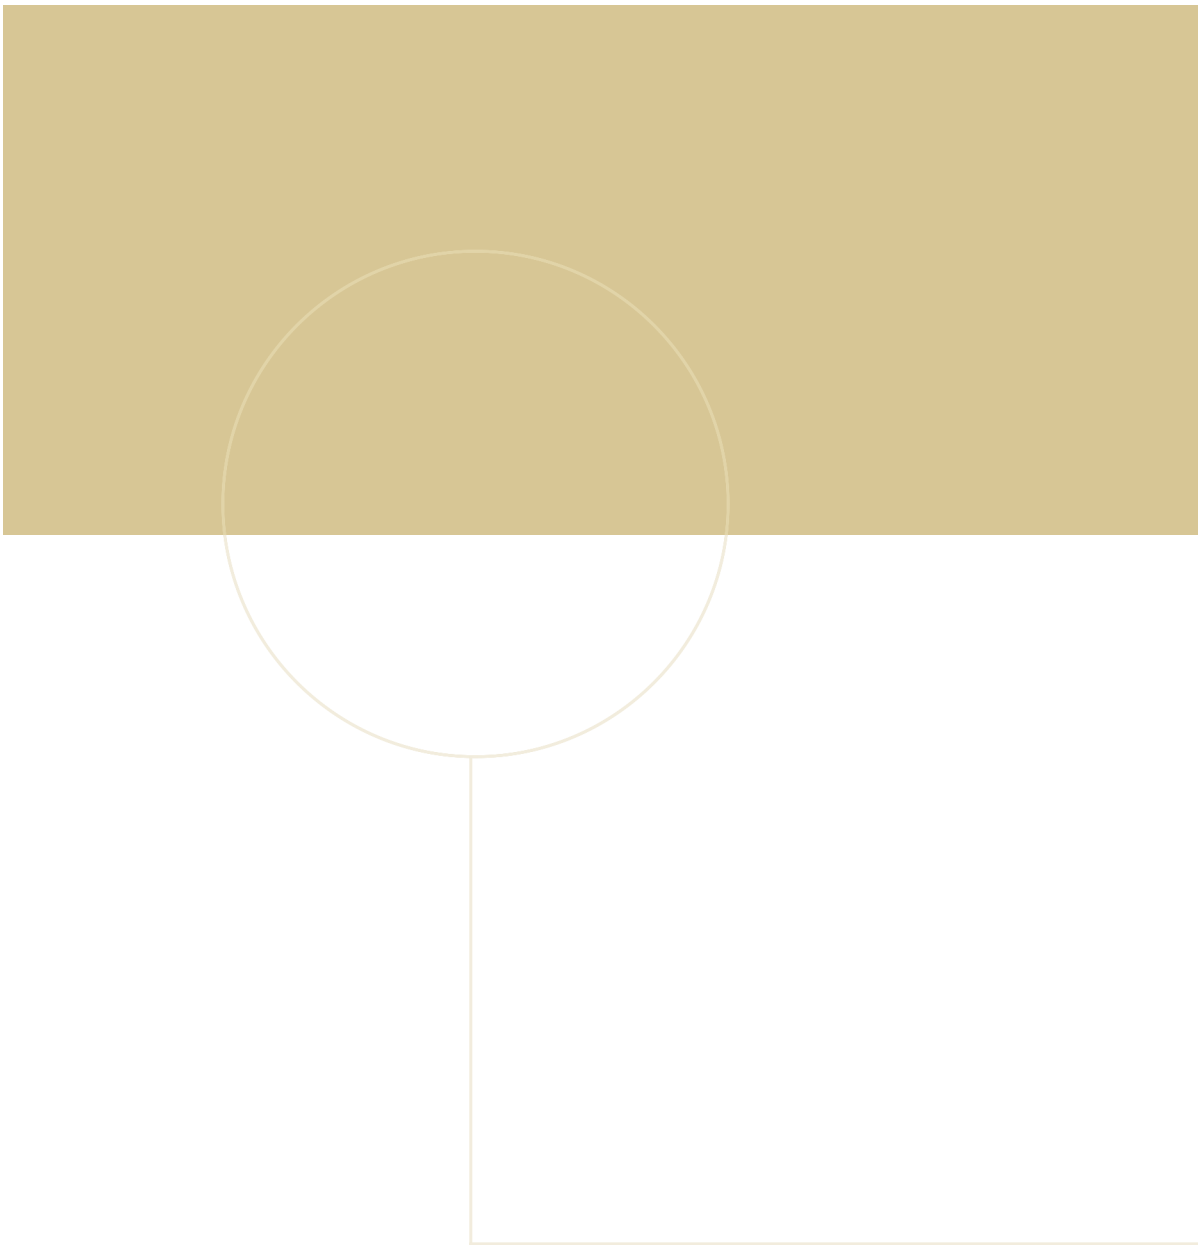

NTNU

Kunnskap for en bedre verden