Concepts for a DLT-based platform for the power generation & trading sector

Alexandre Juncker, Alpiq AG, March 2021

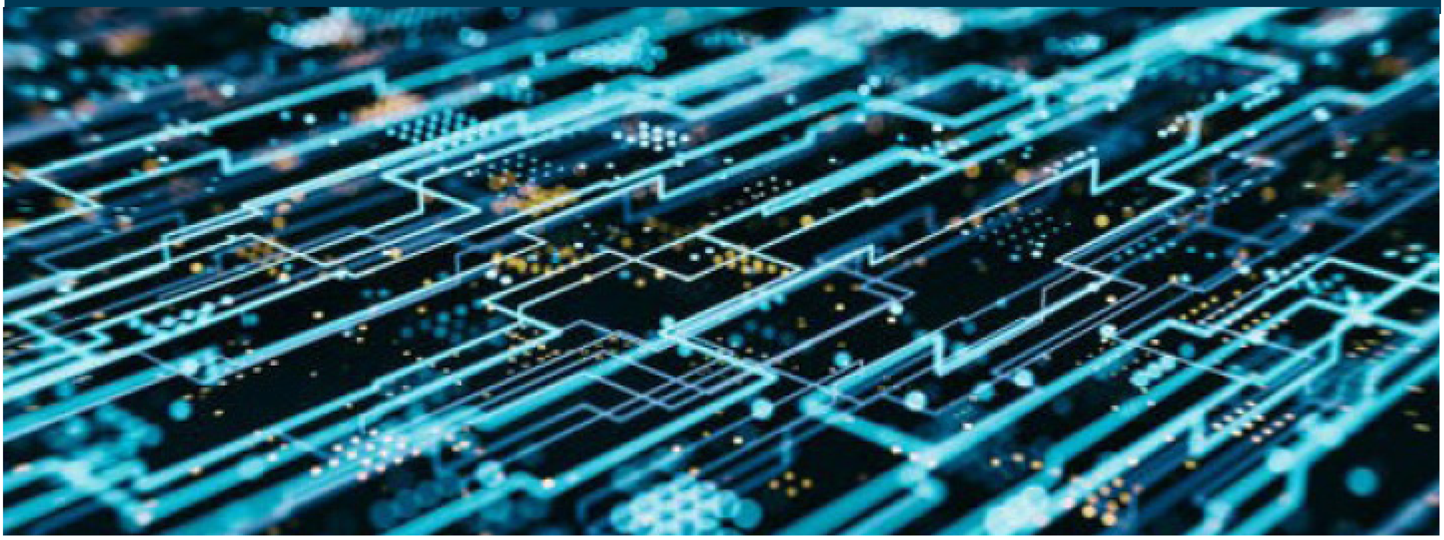

Contents

1 Introduction

1.1 Objective

We consider that the value proposition of distributed ledger technologies (DLTs) as a base layer for the information systems of any given sector is to provide a new, alternative accounting system that promises to be more productive, reliable and efficient than any current system.

Hence, the objective of this working paper is to propose a conceptual design for a DLT-based platform that is capable of supporting the power generation and trading sector.

1.2 Motivation

The current electricity sector is characterised by a range of functions and opportunities, including:

- An increasing consumer demand for new or specialised energy products. This can be traced back to the emergence of environmental awareness and the associated call for (real-time) traceability of energy sources, in particular in terms of the carbon footprint;
- An increasing demand and growing enthusiasm for peer-to-peer energy trading, typically at the level of the neighbourhood in a manner that enhances local sustainability and avoids grid fees; this has been triggered by the expansion of decentralised, small-scale wind and solar production;
- A need to rethink and enhance the management of flexibility, which is due to the massive expansion of intermittent (wind and solar) energies in the network, which calls for an ever more dynamic and efficient electricity market that allows reserve assets to be remunerated in a fair manner and ensures grid balancing at all times.
- A growing interest in driving forward automation, especially for back-office processes, which includes billing and settlements.

DLTs exhibit a number of interesting features that make them ideally suited for the support of both existing and new functions. This is due to the fact that they are, fundamentally, a commonly agreed tamper-proof repository of data. As such, they offer:

- The possibility of built-in traceability
- Automation of asset transfer (and information handling)
- Auditability by design

Therefore, we should use DLTs to design an innovative accounting system to handle, streamline and automate business interactions between players in the electricity market. We believe that we are in an ideal position to apply this technology to address new functions and opportunities.

Since all use cases that have been researched to date (peer-to-peer trading, traceability, grid balancing, etc.) are highly complex, it is obvious to us that a DLT-based platform to support the power generation and trading sector must be designed from the onset to be able to serve all the segments along the value chain, to cover all use cases, including but not limited to peer-to-peer energy trading, wholesale trading, energy traceability, grid balancing, integrated flexibility valuation, etc.

Since such a holistic approach does not currently exist, we included the most fundamental principles and proposals for the design of such a platform in this working paper.

1.3 Preliminary comments

1.3.1 DLT usage is an hypothesis

It is important to note that in the following, we take the hypothesis by definition, insofar as that we explore how to use DLTs to support our business – *all* business. The question whether this is a good idea, and better or worse than existing alternative solutions/technologies is left aside – taking the approach: "If it is in any way possible to use DLTs, let us examine how best to implement this technology."

In this working paper, we will propose various technological options and advance with them in mind. This being said, we will attempt to leave the actual choice of distributed ledger platform open as long as the described design of the system does not require more detailed hypotheses.

Also, the concepts are detailed in a way that explores one way to achieve the goal, which does not mean that other approaches would not work as well, or work better; the objective of this working paper is to fuel research in the sector by showcasing one direction of interest, which can be challenged, refined and reinvented until hopefully the sector converges on a suitable platform design.

1.3.2 Definition of a standard protocol vs. additional applications to be proposed

A DLT-based power management and accounting system must consist of several layers:

- An infrastructure layer that must be based on a standardised protocol agreed within the industry as a whole, or at least by the main players that plan to use it to communicate among each other.
- A set of standard applications providing a range of basic functions that are available to all market participants.
- A commercial application layer where business players can perform their day-to-day operations while relying on standard applications and designing additional applications in order to propose services to customers in the form of instantiated smart contracts that are open to everyone on the basis of a subscription.
- External business processes and tools are interfaced with these applications to facilitate operational realisation within organisations.

In the following, we will propose a conceptual design for the first two layers (infrastructure and standard applications), which first requires the definition of a desired process that must then be supported by the system.

1.3.3 Realisation of the system

Developing a viable version of the concept described in this paper will require the adaptation and deployment of a DLT, the definition of its permissioned nodes, the instantiation of a set of smart contracts, and the development, testing and maintenance of dedicated webpages that retrieve data from the distributed ledger, display it and allow users to send back information/orders to the relevant smart contracts, in order to allow participation in the processes.

We will assume that we will use a permissioned, scaling version of Ethereum (or EOS, or other solutions from this blockchain family), relying on sharding to process information at the various levels of the physical network. In the following, our design will be based on this type of ecosystem.

Whatever platform is chosen, connected smart-meters will have to be installed and declared in the blockchain in order to feed in the generation and consumption data.

ALPIQ

2 General considerations, choices and hypotheses

2.1 Governance

Today, the participants in the electricity generation, distribution and consumption network are readily identifiable, use systematic contractualisation to interact with the other participants, and are accountable for their actions. In other words, it is already a completely interlinked eco-system; and this will remain the case when a DLT-based concept has been introduced.

The consequence is that the network is likely to be permissioned – or at least controlled by its participants – and a solution to record responsibilities must be devised, primarily in order to regulate participants joining or leaving the network.

By default (the most simple approach), we envision an entity that is responsible for this. This entity would be remunerated for performing this task. This could be changed at a later stage, provided that suitable potential alternatives for on-chain governance emerge and prove their robustness.

This implies that the management of this authority should be possible within the eco-system, if only to allow the responsible entity to be changed, or to be revoked by stakeholders; for example, it is conceivable that the stakeholders could vote on decisions. Thought should also be dedicated to fall-back scenarios for the event that the entity defaults.

In the minimum viable solution, it will be supposed that the responsible authority is reliable, stable and not challengeable.

We propose that the entity responsible for controlling participants joining the network is also the grid operator and in charge of balancing the network. In the following, we will suppose that this is the case.

2.2 Discretisation

Two kind of discretisation are defined:

- Discretisation of time allowing all players to refer to the same elementary periods. This could for example be one minute [currently, the shortest time period managed by systems is 15 minutes].
- Discretisation of energy amounts, so that all amounts are a multiple of an elementary amount of energy, for example one Wh.

The discretisation chosen has to be sufficiently small in order to make detailed accounting possible (otherwise we would fail to leverage the a priori opportunities of DLTs), but also sufficiently large that physical measurements still make sense and that no possibility for infinite disaggregation of data is opened up.

2.3 Tokenisation of electrical energy futures contracts

The basic element around which the flows of information in the current system are organised (and potentially even more so in the future) is the contract for the future delivery of electrical energy.

This means that we can either define the recording and follow-up of these contracts as entries in distributed tables that change attributes, owner, etc. – or as non-fungible tokens; in essence the result remains the same.

These "energy tokens" store and convey information throughout their journey; such characteristics may be considered as metadata, and could be standardised as follows:

- Serial number of the token
- Identification of the delivery period concerned:
    - First elementary period of the contract
    - Last elementary period of the contract

- The quantity of the elementary amount of energy is represented in the token, which is understood as a base constant load delivered throughout the time period.
- Identification of the DLT address that will ultimately provide information on the actual energy delivered (entered on the physical network in the event of actual delivery).
    - Note that this can be an IoT address of a smart meter or the address of the trader who decides to open a theoretical position.
    - If this field is specified, the data must be countersigned using the cryptographic signature of the address in question.
    - An attribution table containing the energy delivery address must be mapped to legal entities. A form of identity management on the chain would enable market participants to identify and validate the identity of their counterparties.
    - The type of energy source (wind, solar, hydro, nuclear, coal, gas, etc.) should also be inferable from this smart-meter identification, in the same way as the metadata in the table of the energy delivery addresses.
    - Some form of geographical identification (latitude and longitude) must also be made available in the same way.
    - In the event that the token is not created by a physical producer (and rather by an intermediary, typically a trader), the identification still matters, and would be implemented using a type of "virtual smart meter", which is attributed to the type of energy traded by a given intermediary (i.e. one identification per type of energy traded by each intermediary).
- Identification of the smart meter that will ultimately measure the delivery of the energy to the network
    - The same remarks as above apply: the consumer (legal entity) should be inferable from the identification of the consumption smart meter, and cryptographically verified, etc.
- Price to be paid for the energy upon delivery (also referred to as "face value") – and the currency.
- Price of the token itself. This price is different than the face value to be paid after delivery for the energy itself, and could fluctuate strongly throughout the time preceding the

delivery period. Both the producer and consumer side of the token can be traded independently:

  - Producer position selling price [if void, the token is simply not for sale]
  - Currency for producer selling price
  - Consumer position selling price [if void, token is simply not for sale]
  - Currency for consumer selling price

- Potentially, various standardised contractual parameters [but we expect these to remain very uniform, at least in the wholesale market, because a lack of these would affect the ease of handling and liquidity by multiplication of token types]:
  - Agreed term of payment in days, or rather, the settlement date
  - Penalties in the event of default of production / default of consumption
  - Applicable legal framework
  - Etc.
- Potentially a selection of statuses.
- Environmental impact data:
  - Quantity of CO_2 involved in the production of the amount of energy [note: over the entire lifecycle of the asset, i.e. including construction and projected decommissioning]. This parameter should be established by an independent, recognised audit.
  - [Potentially other types of environmental impact ratings.]
- Information on the possible usage for redispatching and for balancing needs (void if not specified by the creator) [this requires qualification from the balancing authority, either for the participant or for the technology].
  - Mobilisation time for redispatching order
  - Price of redispatching execution
  - Mobilisation time for primary adjustment
  - Price of primary adjustment if purchased by grid control entity
  - Mobilisation time for secondary adjustment
  - Price of secondary adjustment if purchased by grid control entity
  - Mobilisation time for tertiary adjustment

ALPIQ

    - Price of tertiary adjustment if purchased by grid control entity
- Defaulting fees
    - Delay for acceptability of defaulting by the consumer
    - Fee for a consumer default
    - Delay for acceptability of defaulting by the producer
    - Fee for a producer default
- Grid fees (this data is calculated and attached to the token upon delivery), ultimately to be paid by the consumer
    - Transport fee to the grid operator
    - Participation in "total cost of ownership" + fee to the grid owner

Other token characteristics could be added – the only consideration in favour of limiting the amount of information attached to a token is of course the size of the data that will be processed within the DLT and outside of the DLT – which is still a crucial concern with regard to scalability.

2.4 Standardisation of energy categories

When considering the problem of ensuring the traceability of energy throughout the electricity infrastructure, the first consideration that often comes to mind relates to the technology or fuel used for production: hydro, coal, methane, wind, etc.

However, the parameter that actually matters in order to protect the climate is the greenhouse gas emitted – over the entire lifecycle of the plant. So, if we can establish the traceability regarding the accounted CO_2-equivalent emitted per kWh produced at a given plant, we cover environmental concerns and create tools to track the effective CO_2 content of the energy all the way through to the consumer, which could then be used by authorities to promote emissions reductions by implementing incentives or penalties at the appropriate levels, in addition to or replacing existing carbon markets such as the ETS (EU's Emissions Trading System).

However, differentiating energy categories has consequences. These include (1) the concern that if we want to separate different types of electricity, a separate market has to be created for each type and it is highly desirable to have efficient and liquid wholesale markets; (2) the grid would have to balance each energy category in real time. Hence, we cannot design a system where the segmentation is too fine.

From there, it appears impractical to realise markets according to CO_2 intensity. All tokens must be associated with their precise CO_2 content, any incentives or taxes should be implemented as a smart overlayer, the manner of which is yet to be determined. And the wholesale markets should be segmented into 5 or 6 main categories:

- Solar
- Wind
- Coal
- Methane
- Nuclear
- Hydro
- Various biomass
- Others or undetermined bulk

For a power plant to obtain the right to bear the label of a given category, an audit would be required.

Here, however, lies a problem that requires solving: While we want to achieve precise CO_2 accountability in order to drive producers to strive to reduce their CO_2 emissions, we have to harmonise the CO_2 content of coal / nuclear / solar, etc. for the wholesale market. In order to do this, we would introduce a CO_2 authority, and propose the use of ad-hoc tokens with zero energy content and containing a certain amount of CO_2 for players to buy or sell, before or after merging them with their own electricity in order to end up with tokens that contain a standardised amount of CO_2 corresponding to the previous period's average as established by the CO_2 authority. Ultimately, the total CO_2 amount of all tokens that an individual consumes add up and will trigger a proportional payment. Subsequently, the CO_2 authority has to articulate / comply with wider CO_2 markets such as the EU's ETS.

2.5 Geographical scope

Due to the interconnectivity of electricity networks, the ultimate deployment of a DLT-based system only makes sense if it can be adopted at least on a continent-wide scale.

Naturally, however, a potential transition from the status quo to a new infrastructures would be a challenge in itself. Initially, we would consider the implementation in an isolated country or small dependency, for example Andorra, Faeroe, St Pierre and Miquelon, or a Swiss region.

Importantly, introducing the system in a confined geographical configuration does not actually limit the outreach of the design. Indeed, the interface of such a geographically isolated system with the external world could be modelled simply by considering the connecting point as a virtual producer/consumer with the corresponding capacity added to the existing network.

2.6 System architecture

Ultimately, it should be possible to generalise the existing conceptual design to various communities, and then build communities of communities on top. Or inversely, one could consider each participant in the system as a lower-level ecosystem. In this perspective, we propose a design whose principles can be valid at any arbitrary level, in a fractal manner.

Further, an interesting possibility to consider is the matching of a DLT platform sharding structure with the physical architecture of the electrical transmission network. Each level of distribution could correspond to a shard of a larger network, with the main mother chain corresponding to the highest transportation grid. Then, down to the low voltage level, we could associate a shard to each level of distribution, with players thriving at each level, recreating an entire environment of economic activities.

The typical implementation scheme would be to deploy the proposed system at the level of several dozen households, then with another set of households, and subsequently link both with a superordinate layer prior to including more communities, and hopefully building a national level at a later stage.

2.7 Providing tokenised fiat currency in the system

In order to support the automatic settlement of energy transactions (and at a later stage, the payment of all kinds of services provided within the ecosystem), a convenient vehicle of value should be introduced so that it is available within the system and recognised by all market participants.

In the case of a permissionless blockchain, this could be a cryptocurrency that is native to the platform. However, since we envisage the use of a permissioned consortium infrastructure, this option does not exist for us (not to mention that the volatility of such assets is very high, making them unsuitable, at least for the time being).

Therefore, we would have to introduce a currency into the system. The obvious way of doing this is to introduce a tokenised fiat currency on-chain. The official central bank currency of the region where the system operates should be chosen, and a bank entity of sorts would have to be created / appointed that would take over the role of issuing fungible tokens representing one-for-one a unit of fiat. For each issued fiat token, the entity bank would be obligated to deposit the corresponding amount of currency in its vault. Any token holders that comes to the desk of the entity would be entitled to exchange their token for hard fiat currency.

One or several such bank entities could be permitted to provide this service, and would be remunerated for doing so. The actual remuneration model is yet to be defined; it could be one of the following:

- A percentage fee payable to the entity on each transaction involving the tokenised fiat
- A fixed fee per transaction using the tokenised fiat

- [Other models are also conceivable]

In the following, we will not touch on this feature again, but it is understood that most of the smart contracts will require the participants to transfer some sort of tokenised fiat currency to their "account on the smart contract" upfront to enable the instantiated logic to control the corresponding balance and use it to automatically settle the transactions that are initiated by the user.

As an alternative or in addition, some type of mechanism of debt creation within the ecosystem is also imaginable; a type of "I owe you" that would not be collateralised and that would be settled directly between the two transaction parties.

2.8 Regulatory compliance

2.8.1 Macroprudential regulations

Macroprudential regulations such as MiFID, REMIT, FMIA apply in every market.

In the following, we will not focus on this matter, although we would like to note that some sort of supervision will have to be introduced in the system in order to prevent market players from entering into situations that would break these regulations.

2.8.2 Data privacy – GDPR compliance

In a process such as the one envisioned, the data relating to the consumption and generation of energy can be qualified as personal data.

A typical manifestation of the issue is that if electricity consumption data is not properly protected, it could be possible, for example, for thieves to use consumption data of households to plan their break-ins.

In the following, we will examine each of the data protection principles that apply within the GDPR framework.

2.8.3 Lawfulness, fairness and transparency: data minimisation – purpose limitation – accuracy – integrity – accountability

It makes no difference whether a decentralised or a centralised application is implemented. The distributed applications on the platform should be designed to ensure the principles of lawfulness, fairness and transparency are adhered to: data minimisation; purpose limitation; accuracy; integrity; and accountability. A priori, the utilisation of DLT is favourable because it provides a compliant environment that respects these principles; they simply have to be defined as a requirement for the designers and developers.

2.8.4 Storage limitation

What is entered on a blockchain, be it a public or permissioned one, remains on the blockchain forever. Even if the permissioned blockchain might be slightly more favourable in terms of limiting access to the stored data, the only way to comply with data storage limitation requirements is to actually prevent any personal data from ever entering the distributed ledger in the first place.

Thus, it appears that the best way to comply with this requirement is to design the system in such a manner that sensitive information can be concealed. Specifically, we should consider breaking the link to the identity of the individual person in order to comply with this requirement.

Despite the desire for absolute traceability (even if it would be technically possible), we should refrain from systematically handling personal data – in particular the handling of tokens of individual producers and consumers. Importantly, this does not apply to the data of companies; so the requirement is fulfilled provided that a company handles the final aggregation of demand, and the final supply tokens are implemented in an manner that prevents personal data from being recorded (in particular, by applying an off-chain management solution).

This should not prevent peer-to-peer trading, provided it is an explicit choice of the participants to irrevocably deposit their data on a distributed ledger. Outside of these voluntary exchanges, the traceability would be limited to certain generic elements, including limited geographical information, originating at the level of the aggregator.

2.8.5 Confidentiality

This is the most serious regulatory compliance concern that must be addressed.

The general issue is that on public DLTs, data that is considered as sensitive cannot be entered in an openly available manner, as it would be visible to anyone who wishes to join the network – not to mention that it would remain on the ledger forever (see 2.8.4 Storage limitation).

As we plan to rely on permissioned DLT, this problem applies only partially. By controlling who accesses the platform, and each platform being permissioned in itself, obviously the data is not openly available to anyone. However, simply claiming limited access is hardly sufficient, as neighbours on the blockchain could typically directly access the data on the blockchain, which would not comply with data privacy requirements.

Hence, we arrive at the same conclusion as for storage limitation: that compliance with data confidentiality can only be achieved if the sensitive data can be concealed altogether.

The following elements could contribute towards the design of an acceptable (compliant) system:

- To allow peer-to-peer trading, sharing identity is not paramount, but within a potentially small network, associating a pseudonymous account is not sufficient because a certain degree of profiling will be possible, allowing the personal data associated with such accounts to be linked to individuals. Of course, an intermediary (potentially a bot) to link buyers and sellers only after they have come to an agreement is possible, but not fully satisfactory. And since joining the system provides a benefit to the users, it can be assumed that they would be willing to share some of their data in exchange for the benefit. Joining the system is voluntary and sharing certain data is not likely to deter people.

- Similarly, the aggregation of consumption and generation data by a retailer is not possible without passing on a certain amount of data about the community, and thus about the individuals in it. However, in order for the system to work at all, this information is necessary, even today, otherwise, it would not be possible to organise the supply of electricity.
- Tokens created by companies must not contain any personal data about their employees.
- It would be possible to authorise notaries or market authorities to establish a repository of data of the market participants, just as it is compulsory for financial institutions to keep data of their customers, be it merely to enable KYC processes.

Apart from this reasoning, alternatives that should be explored include trusted execution environments, applications that are based on zero-knowledge proof, or also using DLTs for the storage of hashes of information to ensure time-stamped data integrity, at least for certain segments of the value chain.

ALPIQ

3 Description of the process

3.1 Emission

Any player in the power generation and trading market should be entitled to issue energy tokens.

Naturally, in order to find counterparties and convince them to do business, to rely on on-chain exchanges and to communicate with the grid operator and implement balancing in the event of physical delivery, the player will need to be identified and pass certain relevant qualifications by market service providers and grid management entities. Hence every market player will have to be checked prior to actually participating in the market; the actual manner in which this will be achieved is outside the scope of this paper.

When issuing a token, market players will enter their own address both as the generation and consumption location. Producers will use the market to find a counterparty to function as the consumer, and consumers will utilise the market to find a producer. Once a token has been issued, most of its characteristics (amount, delivery time, face value, CO_2 content) cannot be changed anymore; but the consumer and producer identification will continue to be modifiable depending on trading.

At this stage, it does not matter if the energy is physical or not, which means that traders and other intermediaries are entitled to create positions that they commit to covering at a later stage prior to the defined delivery date by acquiring energy tokens with the required characteristics. In the following, we will first examine the communalities of handling physical and commercial energy trading prior to touching on any differences.

Logically, we expect that players will issue tokens in sufficient time for counterparties to have time to match them. However, a high degree of variability in this respect must be expected, for instance with regard to intermittent sources of energy that depend on the weather that can only be reliably forecast a few days or even hours prior to delivery.

In concrete terms, the process of creating a token consists of using an ad-hoc web interface to send a cryptographically authenticated message to the relevant smart contracts, which create the associated token with the desired properties, owned by its creator.

3.2 Trading

3.2.1 Trading principles

Once created, the energy token enters the market to be traded freely. The token has two sides: the producer side and the consumer side. Each can be traded independently.

Traders who wish to sell the consumer or the producer position offer it on the market at a price of their choice. This price can be positive, negative or zero, since it is a future contract for delivery. The price of the collateral energy is to be paid after delivery, but the market value of the position as producer or consumer of the energy can fluctuate. In the basic case, the face value will be the expected price, and the price of the token's producer and consumer sides would be zero. Selling a consumer side equates to selling energy, and selling a producer side means buying energy.

By default, a given side of a contract token has no price (the price metadata field is set to "void"), which means that its owner has not released it for sale. The owner of a token (either freshly created by the owner or purchased on the market) can send information to the system as a message to the token instantiation smart contract to update the price of that side of the token to a non-void value. This identifies the token as offered on the market.

Consumers or retailers (in the sense of demand aggregators) are free to buy the consumer side of tokens from any provider (producer or intermediary) in order to cover their forecast consumption, taking into account the price they are prepared to pay, the source of the energy and any other criteria that are of relevance to them.

Traders, like other intermediaries, are free to participate in the market, improving its liquidity and efficiency while they attempt to make a profit.

In concrete terms, traders or other purchasers of tokens will use an application to connect to the network, explore the content of the distributed ledger, extract/identify the tokens available for sale on each side, display them and perform the typical trading functions, i.e. acquire at a defined price, and place bids in the form of newly created tokens. Here, the design of the user interface / user experience will play a key role.

Whenever a deal is concluded – i.e. when a market player accepts a bid – the smart contract checks the tokenised fiat balance of the purchaser, and whenever sufficient, the settlement of the trade is automatically performed: the tokenised fiat transaction is carried out and the owner of the token (producer or consumer side) is updated.

It is important to note that the change of producer or consumer identification of a token side must be signed both by the previous and the new owner.

Upon change of ownership of one side of a token, the price of the token side is automatically updated by the system to void.

3.2.2 Aggregation and disaggregation of energy products

During the process of trading and day-to-day handling of energy transactions, there is sometimes a need to either merge (aggregate) energy tokens, or conversely to split (disaggregate) large lumps of energy, both in terms of volume and in terms of delivery period. In order achieve this, the approach we envisage is that the market participants retain their position on the one side, and emit new token(s) with the cumulative characteristics (or the split characteristics) on the other side. The following rules would apply:

- CO_2 emissions should be perfectly transposed, as should energy volumes.
- The type of energy must be homogeneous.
- The origin must also be homogeneous from a single geographical shard.

3.2.3 Marketplace(s)

One or several centralised or decentralised marketplaces are conceivable, with associated front-end software. The platform(s) would feature both simple and aggregated products, allowing buyers and sellers to find each other.

Settlements on the market places could occur off-chain, but in order to leverage the full potential of DLTs, the best practice is to implement the direct settlement of transactions on-chain. Accounts in fiat currency also managed by the token-managed smart contract would thus enable the instantaneous payment for token sides between counterparties.

In other words, in this decentralised process design, the marketplace is merely a means to explore the of statuses of available tokens, the display of the results to the market players, and the transmission of buy/sell orders and the subsequent change of token ownership in exchange for tokenised fiat currency, settled on-chain.

The question whether a situation with a number of competing marketplaces, especially with regard to liquidity, will arise, or if we will end up with a single public service exchange remains to be seen, but this will have little impact on the rest of the distributed energy management processes.

3.2.4 Generation and consumption incidents

Incidents can always occur in the energy sector, for example unscheduled downtime, strikes, etc. In the event the producer or consumer of a token wishes to withdraw from the contractual obligation, they would have to negotiate this with their counterparty; this could even be included as a feature of the token, for example with a pre-agreed fee. This would offer both parties planning security for the event that a contract for delivery is cancelled.

If an agreement cannot be found, the defaulting party would have to offset their inability to deliver the sold electricity or to accept the delivery of the purchased electricity on the market.

3.3 Pre-delivery, summary of positions and freezing

3.3.1 Balance of positions by design

When the date of delivery approaches, it must be ensured that generation and consumption commitments balance during the period of time concerned.

As long as the grid operator can access all the outstanding tokens emitted by all market players, balancing can be ensured: Each token has an associated producer and a consumer – tokens with identical producer and consumer can simply be discarded.

The only risk that has to be managed is that market players might take positions that they are unable to fulfil. The approach to mitigate this risk consists of several measures: financial penalties would be implemented on any discrepancies that occur once the delivery period has passed; limitations of open positions could be enforced as a safety mechanism, with the grid authority deciding which limit a given market player should be granted depending on their circumstances; dynamic provision of reserve energy by the grid operator could be implemented depending on the observed "dangerous" positions of market players (the cost of which would, again, be passed on to market players by means of defaulting penalties).

This being said, it would remain the responsibility of each player to ensure they ultimately own the positions that correspond to their requirements – for example, pure traders should cover all their sales by taking up the same volume of purchases.

3.3.2 Procurement of tokens by the grid operator balancing authority

Measures for grid stability adjustment

The grid operator must have access to a generation capacity margin in order to fulfil its role of balancing production and consumption. Conventional primary, secondary and tertiary control energy must be provisioned because more or less severe unbalances cannot be ruled out due to deviations by producers (accidents, weather conditions etc.) or consumers (forecast errors, exceptional events, etc.).

Hence, the grid operator should have the capability to procure energy tokens with a metadata identification as "for balancing purposes". The face value of such energy tokens would not be so important – rather a contract price would remunerate the generation capacity for its availability.

The operational activation of the capacity could also be performed on-chain. The actual mechanism to communicate the availability of primary, secondary and tertiary flexibility to the grid in real time is beyond the scope of this working paper, but it should be kept in mind that DLTs could also be used to register the operator's grid stability commands and automate the response by industrial plants.

Measures for grid losses

The grid operator must also purchase the consumer sides of tokens in order to "feed" them into to the transmission or distribution grid to compensate for the Joule effect that occurs during electricity transmission. As all the accounting of energy is carried out within the system, these grid transmission losses must be factored in in advance, just as is the case in the current system.

The manner in which the associated costs are to be passed on to the end users is to be established separately.

3.3.3 Positions freeze and redispatching

In principle, there would no need for a positions freeze function, but for practical reasons, it makes sense to establish one, as it would allow the grid operator to take measures for redispatching and balancing by considering the positions taken by the market players.

Prior to the delivery date, the grid operator must ensure redispatching. This involves observing the net energy flow between the players, evaluating the risk of grid congestion and taking measures to balance the various parts of the grid where necessary.

In practice, these measures will be recorded as energy token sides that are modified (reallocated) through the change of producer or consumer identification, with financial compensation. The details of this process are outside the scope of this paper.

3.4 Physical delivery

The electricity is physically delivered during the elementary period of time specified on the token.

The connected smart meters measure each producer's production of energy: They are identified cryptographically on the DLT network, which allows them to push data as oracles. Similarly, each consumer's consumption is measured, and the data is pushed onto the DLT network.

The data is generated by the individual smart meters and is registered in a suitable on–chain table.

3.5 Post-delivery accounting

3.5.1 Discrepancy reconciliation

Since actual deliveries always deviate at least slightly from the forecasted positions, deviations must be accounted for.

Demand and supply must always be balanced in real time, not only in total, but also by type of energy. Even though the physical networks used are the same, the precise real-time adjustment of production would be performed independently for wind, nuclear, coal, etc.

The following four cases have to be dealt with when comparing what the smart meters reported with the sum of the existing tokens (this information is accessible by retrieving the measurements of each smart meter, and calculating the delta compared to the owned and created tokens allocated to this meter):

- More energy is produced on the producer side than scheduled: An offset token is created automatically, with a face value corresponding to the weighted average of the electrical energy price observed in the elementary period minus an n% discount. The consumer of this token is defined as the balancing authority.
- Less energy is produced on the producer side than scheduled: A token is created, with the consumer as the incorrect player, and the producer as the balancing authority; the token face value is the average price plus an n% penalty.
- Less energy is consumed on the consumer side than scheduled: The system automatically creates an energy token that makes up the difference; the declared producer is the customer, and the declared consumer is the balancing authority. The face value is the average price of the period minus n%.
- More energy is consumed on the consumer side than scheduled: The consumer is declared the consumer of a token that is automatically created to make up the difference, with the producer being declared as the balancing authority and with face value as the average price plus n%.

Once these four cases have been resolved, all energy is accounted for. The balancing authority has levied funds that can be allocated for the provision of its control capacities, and all the tokens created by consumers equate to those created by producers. All the high-precision information about consumption is available simply by exploring the token characteristics.

Note: Certain jurisdictions allow for “day-after” markets, which are also an interesting way to perform (a part of) the reconciliation. What matters is to incentivise the market players to facilitate the task of the grid balancing authority. A thorough study would be required to design an alternative process based on day-after markets, which would take effect prior to the automatic reconciliation of positions.

3.5.2 Grid transmission and distribution remuneration

With our system, we completely separate the transmission costs from the energy costs. We will also attempt to decompose this transmission and distribution price in order to maximise the informative value.

Balancing service

It is only logical for the service of balancing the network to be remunerated using the penalty fees collected for the reallocation of energy among market participants. The details of how this accountability should be engineered goes beyond the scope of this paper, but the relevant point is that whatever the solution chosen, it will be possible to implement and hence automate it using smart contracts.

Electrical energy transmission service

Upon delivery, the grid operator accounts and establishes the transmission costs incurred during the elementary period. In particular, these should cover the Joule effect and the redispatching costs.

The net energy consumed by a given participant in the network can be viewed simply by exploring the volume contained in all the tokens in which that participant is involved over the delivery period under consideration. The allocation of the cost pro rata of the consumption can be performed in a simple, direct and, again, automated manner.

Remuneration for ensuring security of supply

One final element of grid transmission and distribution remuneration is the security of supply. While driving forward clear accountability of grid utilisation is interesting and allows for the incentivisation of local production and consumption, security of supply also has an inherent value, and it would make sense to account for this component separately.

The annualised "total cost of ownership" (+fee) could be passed on to the end users via a suitable mechanism that remains to be established: either a fixed fee per market participant, a pro-rata of their total consumption (counted both locally and non-locally), or other mechanisms that make sense from a social perspective. The underlying logic should be to reduce some of the incentives for decentralisation that would otherwise result in a feed-back loop undermining the sustainability of the interconnection grids – which this working paper deems non-desirable from the perspective of security of supply.

An application could be implemented that automates this remuneration, but since it will depend on the observation of the consumption / production of maximal amounts over a period of time, it should be performed separately from the day-to-day settlement of transported energy tokens.

3.5.3 Realisation of traceability

In the exposed system, we provide traceability by ensuring transitivity of the characteristics of the energy from the producer via any intermediaries right through to the consumer.

However, not all characteristics can be traced in this way, especially the precise geographic origin information would most likely have to be obscured for data protection reasons.

Provided that intermediaries can be convinced to allow access to their data, it would be possible for an energy accounting service to explore the ledger. Using exhaustive records of transactions for a specific end consumer, the origin of the energy could be calculated on a pro rata basis. In general, the result is likely to show a relatively scattered status. Apart from this, the transitivity of most of the data over the course of atomic tokenisation should allow the most important characteristics of the energy to be transmitted to the end customer in a satisfactory manner.

3.6 Settlement

3.6.1 Regular settlement

The final step – the payment – this is quite simple in comparison to all the energy accounting.

Importantly, market participants must have transferred sufficient cryptocurrency funds to the relevant settlement account prior to trading.

The settlement function screens the non-settled tokens that are due. For these tokens, it transfers payment from the participant that owns the token to the participant that created the token, according to the price determined on the token.

Naturally, this also applies to all the intermediaries that hold or have created tokens, and also to the grid operator.

In a similar manner, a specific settlement logic has to be enforced for the flexibility providers that are under contract with the grid operator.

Payments are typically made in the tokenised fiat currency, which is the same for all the players in the ecosystem [at least in the initial phase].

If a payment is unsuccessful due to insufficient balance, one solution could be to postpone the settlement date while adding a fee to the amount to be paid.

Ideally, just as in the conventional system, settlement would be performed by netting positions, so that treasury limits would not be a restrictive factor preventing intermediaries from taking positions.

A part of the settlement would be used to pay a fee for the CO_2 content consumed.

3.6.2 Counterparty default

So far, we have disregarded the possibility of discrepancies in the financial robustness of token purchasers. However, as these are contracts for future deliveries, the identity of the individual purchasing the token actually matters. There are several approaches to deal with the risk of counterparty default.

- Firstly, it would be conceivable to introduce ratings of market participants. Under a certain threshold, a certain degree of upfront collateralisation could be required to mitigate the risk. Depositing funds in escrow in the smart contract would be easy to implement.
- Secondly, an insurance service could be provided on-chain in the ecosystem.
- Thirdly, a straightforward solution would be for the participants to agree to assume the risk as a dedicated intermediary in the market.
- Finally, it would be possible to implement a penalty mechanism: For example upon default, the system could automatically exclude the player from further participation in the system. Many types of penalties could be implemented, for example a partial or total ban, fines, etc.

4 Generalisation of the model to multiple grid levels

In the previous chapter, while examining individual provisions for the generalisation of the system, we laid out how the system could work at the level of a single grid layer. Namely this can apply to a microgrid or to the wholesale market containing only large players.

However, if we want such a system to be rolled out to the entire market, from the top – the continental level – to the bottom – the neighbourhood – we have to extrapolate these principles to an architecture with multiple grid levels. All levels would operate according to the same principles, with the lower levels functioning as shards of the upper levels.

Here, we envision a landscape with a number of physical grid operators (which are most likely usually also the grid owners) that receive remuneration on each level according to the approach described above; while at the upper level, the sub-grids (such as DSOs compared to TSO) can be viewed as individual market players.

When an energy deal occurs between two players that are not in the same physical grid, the problem is that any other grids that need to be traversed must be remunerated. In practice, this will correspond to having a producer that is not in the same shard as the consumer.

In the following, we will examine how an effective management of energy tokens can be achieved, in particular for trading, and investigate the various tasks that grid operators have to perform in their day-to-day operations.

4.1 Trading

If a consumer trades with a producer that is one or several grids apart, the current solution foresees issuing several energy tokens instead of just one.

Namely, a token would be created linking the producer to a dedicated non-personal address within the same shard that has the purpose of interfacing and that is registered on the superordinate shard. In this way, several vertical levels can be bridged, with one token created in each shard to represent the energy flow. Then, the same process is applied downwards until reaching the actual consumer. Now we have energy tokens mapping the complete physical flow throughout the grids.

From this situation, the objective is to allocate the grid usage costs to the market players. Note that the proportion of grid fees charged to each of the players is defined when they conclude the deal, and this split will be used recursively.

4.2 Redispatching

At any level, redispatching can impose a forced reallocation of energy flows. If one token touching the root address has to be cancelled for redispatching reason, then the whole chain of tokens has to be deleted as a consequence.

As this would impact more than just one grid, we absolutely want to avoid this scenario. In order to ensure that this will not happen, each "bridge" should have a maximum capacity that the grid operator of the superordinate level must define. It would then only be possible to buy a token that uses this bridge if the capacities of all bridges linking both players still have sufficient margins.

4.3 Post-delivery reconciliation

The discrepancies between the planned and the actual situation are managed conventionally within each grid at both ends, with no impact on the chain of tokens. On all other levels, the precise flow is consumed and produced.

4.4 Grid fees

This is where things get interesting. At each grid level, the transmission fees are calculated for each period and allocated on a pro-rata basis to the energy tokens that were handled during that period.

The solution are impersonal addresses that ensure that in addition to the price of the token, each interaction between levels on the path to the end customer is subject to a charge for a conservative estimate of the cumulative grid fees. This conservative estimate of the transmission charge should be visible to the end customer and is due at the moment the consumer side of the token is purchased.

One of the consequences of this approach is that only consumers can purchase energy tokens from outside of their shard; otherwise it would require charging this cost to consumers without them having any control over it. However, the purchase of energy from a specific precise origin is only ever likely to occur at the initiative of end consumers.

Another consequence is that the peer-to-peer purchase of energy outside the local network would be heavily de-incentivised.

5 Conclusion

5.1 Summary

This working paper outlines a proposal of how end-to-end power generation and trading processes could be transposed to a DLT-based platform.

Throughout this paper, we have demonstrated that a DLT-based system is extremely well suited to tackle the new functions and challenges facing the electrical energy sector, in particular in terms of traceability and settlement automation. The streamlining of interfaces and the generalisation of principles regardless of the grid level are additional strong points of the design and make it easily applicable to peer-to-peer trading as well as wholesale. Ultimately, it provides faster and more transparent interactions, makes markets more efficient and enables preferred means of production to obtain more financing and flexibility management to be handled more dynamically and automatically.

Many points remain open, in particular because the technology is still in its infancy. However, provided the progress in terms of scalability, data privacy management and governance mechanism continue, the deployment horizon of such a concept could be as close as a couple of years.

5.2 Swiss DLT-for-Power standardisation initiative

Alpiq has driven forward the unification of all the players in the Swiss power generation and trading sector in an effort to join forces to explore the concrete features of a common platform. This “DLT-for-Power” initiative was launched in November 2019 and it is being developed under the umbrella of the Swiss Association for Standardization.

Ultimately, the goal of this initiative is to develop and deploy a suitable infrastructure. Over the short term, expert groups are working to define principles in a number of technical focal points covering:

- Distributed identity management, confidentiality and privacy
  - Definition of the technical principles
  - In particular, examination of the possibility to utilise existing standards
  - Addressing of GDPR regulations constraints with respect to distributed identity management
  - More broadly, definition of principles to ensure data confidentiality in the system
- Governance principles on various levels
  - Consensus participation
  - Distributed logics responsibilities
  - On-chain authorisation governance
  - Grid balancing responsibility
  - Possible physical network ownership scenarios
- Scalability and interoperability
  - Study of the scaling principle, in particular sharding in parallel to the physical architecture
  - Examination of the question: One or several DLTs?
  - Definition of principles addressing the interoperability challenge
- Shaping the execution environment
  - Definition of the principles regarding what information should be on the ledger and what information should not
  - Definition of principles regarding which functions should be executed on-chain versus which functions should be based on simple hash notarisation
  - Definition of the limit between the common-layer infrastructure and the competitive-layer application
  - Exploration of the interest of using contracts for future delivery of energy as the basic element of the system; drafting of the basic characteristics of such a token

- Validation or rejection of the concept of introducing a tokenised fiat currency in the distributed environment

5.3 Alpiq's technological lead

Alpiq is extremely conscious of the relevance of digitalisation in its sector. In addition to its prominent role in the DLT space, the company is striving to monitor, acquire expertise, and proactively explore applications in other fields of digitalisation, including artificial intelligence (AI), the Internet of Things (IoT) and 5G telecommunications.

Bellow, we will briefly outline Alpiq's vision and endeavours in each of these fields.

5.3.1 Artificial intelligence

The opportunity to automate large parts of energy trading using DLTs will open the market to participants that have hitherto been excluded. Individual prosumers and energy communities with renewable assets will be able to trade power among each other and with already established market participants. In addition to market access, these newcomers will also need production and consumption forecasts and automated dispatching for flexible assets to make their participation feasible.

DLT-for-Power is thus expected to create a demand for AI solutions to forecast intermittent power from renewable assets and inflexible demand for prosumers using machine learning, to automate the dispatching of flexible assets, and to optimise the price for each bid and offer. Even established energy producers will find the increase in data overwhelming and will require AI to free their traders from real-time trading, so that they can focus on mid- and long-term strategy development. Alpiq Digital & Commerce offers custom power forecasting that model assets based on machine learning and develops optimisation solutions that save energy, reduce costs, and offer the automated solutions required for new business models.

5.3.2 Internet of Things

Real-time asset monitoring and control possibilities are key elements of DLT use cases. Using the IoT, Alpiq has already connected many different prosumer assets ranging from electric car charging stations to large-scale industrial batteries. Connecting these assets opens up the possibility to offer flexibility for the grid or local energy communities.

Fraction-of-a-second streaming of asset data opens up improved understanding, and together with AI applications, it can support other use cases, such as predictive maintenance, that save costs and reduce downtime.

We believe that the flexibility provided by connecting assets is the key to support and drive forward the change in our industry.

5.3.3 5G

Alpiq is also involved in exploring the opportunities offered by 5G in the power generation and trading business. 5G can be viewed as an enabler for the massive expansion of IoT in many new areas and for a boost of the deployment of robotics.

Ultimately, if blockchain delivers on its promise to reduce transaction costs to just fractions of a cent, the implementation of micro-transactions between robots is likely to require the support of 5G.

Legal disclaimer

This working paper is made available on the express understanding that it will be used solely for general information purposes. The material contained in this working paper should not be construed as relating to accounting, legal, regulatory, tax, research or investment advice and is not intended to take into account any specific or general investment objectives. The material contained in this working paper does not constitute a recommendation to take any action. Before acting on anything contained in this material, you should consider whether it is suitable to your particular circumstances and, if necessary, seek professional advice. No representation or warranty is given that the material in this working paper is accurate, complete or up-to-date.

The material in this working paper is based on information that we consider correct and any statements, opinions, conclusions or recommendations contained therein are honestly and reasonably held or made at the time of publication. Any opinions expressed are our current opinions as of the date of the publication of the working paper only and may change without notice. Any views expressed in the working paper represent the views of Alpiq only, unless otherwise expressly noted.

While the material contained in this working paper has been prepared in good faith, neither Alpiq nor any of its agents, representatives, advisers, affiliates, directors, officers or employees accept any responsibility for or make any representation or warranty (either express or implied) as to the truth, accuracy, reliability or completeness of the information contained in this working paper or any other information made available in connection with this working paper. Neither Alpiq nor any of its agents, representatives, advisers, affiliates, directors, officers and employees undertake any obligation to provide users of the working paper with additional information or to update the information contained therein or to correct any inaccuracies which may become apparent – except through its own revision procedure. To the maximum extent permitted by law any responsibility or liability for this working paper or any related material is expressly disclaimed provided that nothing in this disclaimer shall exclude any liability for, or any remedy in respect of, fraud or fraudulent misrepresentation. Any disputes, claims or proceedings in connection with or arising in relation to this working paper will be governed by and construed in accordance with the laws of Switzerland and submitted to the exclusive jurisdiction of the courts of Zurich.