Blockchain-based Local Energy Grids: Advanced Use Cases and Architectural Considerations

Marco Peise, Jörn Kuhlenkamp, Anselm Busse, Jacob Eberhardt, Max-R. Ulbricht and Stefan Tai
Information Systems Engineering
Technische Universität Berlin
Berlin, Germany
{mp,jk,ab,je,mu,st}@ise.tu-berlin.de

Jörg Baus
Deutsches Forschungszentrum für
Künstliche Intelligenz GmbH
Saarbrücken, Germany
baus@dfki.de

Martin Kassebaum
Enercity AG
Hannover, Germany
martin.kassebaum@enercity.de

Thorsten Zörner
STROMDAO GmbH
Mauer, Germany
thorsten.zoerner@stromdao.com

***Abstract*—Energy markets worldwide are undergoing a transformative change from fossil energy sources and centralized setups to renewable energy sources and highly decentralized markets with peer-to-peer trading. Blockchain technology has been considered a fundamental enabler for building decentralized energy systems and services. However, existing blockchain-based energy solutions may not be transferable to new contexts, as regulatory, infrastructural, and economic criteria and constrains may not fit or have not been made explicit with existing work. To this end, we present a 3+1 layer model for use case identification and for architecting their realization to ensure coverage early-on of the different critical dimensions. We report on findings from a major project running in Germany by describing five advanced use cases and discussing important architectural considerations.**

***Index Terms*—Blockchain, Use Cases, Trading Renewable Energy, Architectural Implications**

I. INTRODUCTION

Energy markets worldwide are currently in the middle of a paradigm shift to address climate change. A major element is the transformation from fossil fuels towards renewable energy sources, which is typically coupled with a parallel shift from centralized and large-scale to decentralized and small-scale, even household-based energy generation.

Blockchains have been proposed as a technology to address some of the decentralization challenges that come along with this major transformation, including peer-to-peer energy trade and censorship-resistant, manipulation-safe transaction processing. Various studies, surveys, proposals and concrete projects have been conducted [1]–[4], which, in union, demonstrate not only the potential but also the viability of blockchain-based energy systems.

Taking a closer look, however, most of the existing blockchain-based energy solutions are isolated solutions that are not easily transferable to a new context. Regulatory (national or cross-border), infrastructural, and economical dimensions may discourage or even prohibit adoption of an existing solution but require development of a new solution. Specifics related to electricity tax, grid charges or renewable levy, along with constraints imposed by existing (legacy) energy infrastructure and business models in-place are examples of important properties that require careful attention when engineering blockchain-based energy systems and services, but may not be clear with existing work or were simply ignored.

The German Federal Ministry for Economic Affairs and Energy (BMWi) currently is supporting a consortium consisting of industry and academic research partners working on blockchain-based decentralized energy. The project *BloGPV*[1] explores the use of blockchain technology in combination with energy storage for communities of solar energy prosumers. The objective of the project is to consider all of the dimensions above to enable small-scale prosumers in local grids to operate in economically profitable ways, maximizing the usage of jointly generated energy and minimizing the amount of energy bought from a utility.

In this paper, we report on the challenge related to identifying and implementing advanced blockchain use cases as part of the larger engineering process in *BloGPV*. We present a "3+1 layer model" intended to aid in use case identification and architecting their realization by enforcing consideration and explicit modeling of the different dimensions required. Further, we describe five advanced uses cases for blockchain-based energy trade, namely

- **Tokenization**: Proof of trade-able energy created from renewable sources
- **Accounting**: Privacy-preserving use of smart meter and account balances
- **Netting**: Transparent, law-compliant participation in EEG levy exemption
- **Contracting**: Matching energy demand and supply ex-ante consumption

---

[1] https://www.blogpv.net

- **Compensation**: Automated enforcement of electricity trading contracts

and report on architectural options and decisions taken when engineering the use case solutions.

The five use cases follow the 3+1 layer model proposed so that their origin and description becomes transparent in a way to allow architects and engineers of future blockchain-based energy systems and services to better decide about applicability and transfer of proposed concepts and solutions.

The discussion of major architectural options and decisions needed for each use case, furthermore, demonstrate the significant but non-obvious impact of these on deployment, procurement, and expenses for setup and maintenance that architects and engineers must be well-aware of.

II. Background and Motivation

Blockchains synthesize concepts from distributed computing, applied cryptography, economics and other to enable the building of decentralized, censorship-resistant and manipulation-safe transactional systems [5]. Blockchain-based systems provide for unprecedented trust guarantees and have been found to be beneficial also for the domain of energy distribution. Instead of requiring trust in a single party, like a large energy utility, a network of peers executes programs (smart contracts) and reaches consensus on the results in a trustless manner. In the *BloGPV* project we leverage blockchain technology in two ways:

- to process data in a trustless and censorship-resistant way. Blockchain-based processing protects stakeholders from manipulations and faults.
- to record data in a permanent way by anchoring it in the blockchain's immutable history. Here, the blockchain serves as the single source of truth.

From an applications and value proposition perspective, the goal of *BloGPV* is to ensure economical operation of decentralized, small-scale power generation, for example, by households with solar installations, even beyond government-issued and legally guaranteed feed-in tariffs. The aim is to create targeted incentives that ensure the continued operation of individual generation plants. The vision is to introduce secure and highly automated accounting and billing services for community-driven, peer-to-peer energy trade. To this end, blockchain technology in combination with an intelligent, decentralized battery storage management are needed and to be integrated with existing physical energy infrastructure. At the same time, however, the privacy of prosumer data related to energy production and consumption must be ensured.

The general idea of blockchain-based energy systems and services has been voiced frequently both in industry and the scientific community [1]–[4].

As mentioned in section V, related work on blockchain-based energy systems and services, while manifold, has not reported on a systematic engineering method to identify and document use cases that explain the different regulatory, infrastructural and economic considerations present when building the decentralized solutions. For a given context like the *BloGPV* project, where blockchains and battery storage are used in combination to enable privacy-preserving energy trade, a multidimensional approach to use case engineering has proven to be critical both for communication within the project and with other related efforts.

III. 3+1 Layer Model

We present a "3+1 layer model" for identifying use cases when engineering blockchain-based energy systems and services. The three main layers of the model building on each other are *physical infrastructure*, *trustworthy computational infrastructure*, and *application services*. Further, we define the crosscutting "layer": *market and regulations*.

Layer 0 – Physical Infrastructure: *Layer 0* includes all aspects of the physical energy infrastructure. Here, we focus on physical infrastructure components that can be observed and/or controlled via an API. Common examples of such physical infrastructure components are solar installations, battery storage, consuming units, power lines, and smart meters. A participating household installs a subset of these physical infrastructure components. The public energy grid connects households with each other and the utility.

As an example, Figure 1 illustrates four households participating in energy trading on *Layer 0*. Each household connects to the grid via a grid connection point. All households can withdraw energy from the public grid and consume energy locally. Besides, Households B and D can produce solar energy locally and feed energy back into the public grid. Households C and D have access to battery storage that allows to store and release energy locally. On *Layer 0*, energy trading between two households is perceived as feeding energy in and withdrawing energy from the public grid simultaneously.

Layer 1 – Trustworthy Computational Infrastructure: On *Layer 1*, fundamental functionalities that serve as generic building blocks for applications and services reside. These functionalities make use of observations made on *Layer 0* or control physical infrastructure in *Layer 0*. In general, Layer 1 has to ensure that functionalities are trusted and execute promptly. In order to be accepted by all stakeholders, the system has to be fully trusted. As energy trading will be decentralized, trust also has to be achieved in a decentralized manner. For that purpose, *Layer 1* utilizes blockchain technology. Thus, *Layer 1* is perceivable as a blockchain-driven compute platform. As such, *Layer 1* comprises not only the selection of blockchain technology but also related technology and tooling along with their programming models to develop and operate the blockchain-based energy systems as needed.

For illustration purposes, we extend our example from above (see fig. 1) by a *Layer 1* view. Each household now operates a full Ethereum node. The five nodes form a private and permissioned blockchain network. Due to confidential computing needs, *Layer 1* can use advanced cryptographic techniques, e.g., verifiable off-chain computations based on zero-knowledge proofs [6].

Layer 2 – Applications/Services: *Layer 2* consists of feature rich applications and services that can target all stakeholders

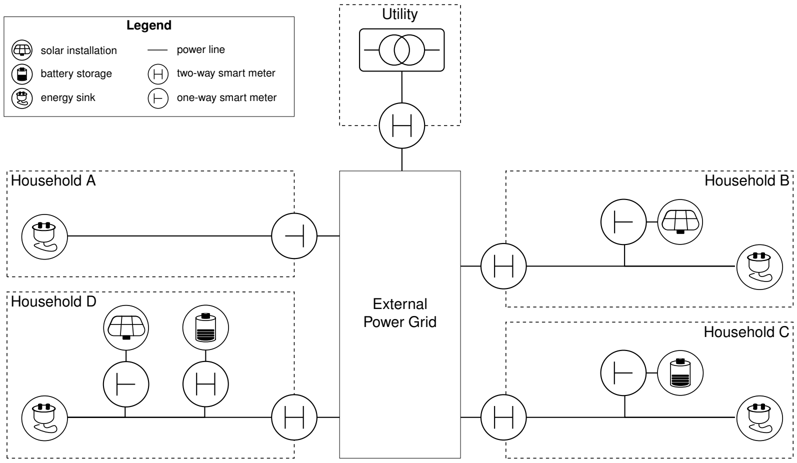

Fig. 1. Illustration of a *Layer 0* view on four households participating in energy trading.

in the energy system including prosumers, utilities, or governmental bodies. Services may include billing, intelligent energy usage analysis and visualization, or data-driven future investment services. The services will build on the fundamental functionalities provided by *Layer 1*. Because of the guarantees provided by *Layer 1*, *Layer 2* can focus on application-level functionality and does not need to address functionality of the lower energy network, though additional services and features beyond that may be devised.

Layer +1 – Market and Regulations: This crosscutting "layer" includes important market and regulatory specifics of the geographic region(s) in which households participating in the energy trading reside. As such, the last layer is orthogonal to the other three layers. This layer is likely to define the concrete business model of any energy trading application implemented on *Layer 2*. Besides, specific regulatory obligations may be stated here that can require the use of a specific blockchain technology on *Layer 1*. For example, the need for mechanisms to ensure privacy-preserving data processing in compliance with the GDPR may be stated here.

IV. USE CASES

As mentioned above, the identification of use cases when engineering blockchain-based energy systems or higher-level services can be challenging concerning the multiple dimensions to be considered. Therefore, we introduced our *3+1 Layer Model* in the last section to support architects and engineers during the process of decision about concepts or solutions to be integrated into newly developed systems.

During the requirements analysis phase of the *BloGPV* project, we identified five critical use cases to derive requirements for the overall system design. Without claiming completeness, these use cases can act as blueprints for the development of future blockchain-based systems for community-centered trading of renewable energy.

Although the use cases identified are very different, we have observed similarities between all of them regarding the various layers from which we derived our model. Since the focus of the project is on community-centred energy trading, it is not that surprising to find the most common requirements in the layers that strongly depend on the location of the actors, namely *Layer 0* that represent the physical infrastructure as well as *Layer +1* as an agent for all regulatory and market-oriented aspects of the particular use case.

Concerning the physical infrastructure, all identified use cases rely on four key components (as depicted in figure 1). Firstly, households acting as members of a community and can take on two different roles, consumer or prosumer, depending on the fact if energy is solely consumed or additionally produced. Secondly, a Utility that can buy and sell energy. Thirdly, a power grid that interconnects households and the utility. And lastly, and most importantly, a reliable and trusted smart metering infrastructure measuring all consumed and produced energy and providing access to these measurements.

Regarding regulatory and market-specific aspects of the identified use cases, it seems comprehensible to assume that if the actors form a local community, they are all located in the same territorial scope of legal regulations and, therefore, the same economic mechanisms and conditions. The most prominent similarities on this layer in *BloGPV* are uniform trading periods of 15 minutes and legal obligations regarding privacy-preserving data processing according to the General Data Protection Regulation (GDPR) or regarding specific regulations for production and trading of renewable energy (RE).

Three of the five use cases have already been implemented as prototypes (*Tokenization*, *Accounting* and *Netting*). In our proof-of-concept implementation architecture [7], specific system components represent each role of our Use Case. For example, the actor smart meter operator is represented by a smart meter. The Household Processing Unit (HPU) system component represents a household and the Netting Entity can be represented by providers.

Based on our experience, we present selected important decisions and their implications for the implementation of systems. In the following sections, we give an extensive description of all five identified use cases that we evaluated during our project's requirement analysis phase. Besides, we discuss various options and their implications as architectural decisions at the end of selected use cases. This is by no means a comprehensive overview of all possible decisions for software architects, but rather an anecdotal summary report of a research project under real conditions with selected examples as abstract decision support.

A. Tokenization

Aim: Proof of trade-able energy amount created from renewable sources

Description: In *Layer 0*, a prosumer with a solar installation can generate solar power and either consume, store (within battery storage), or feed this energy into the external grid. When producing solar energy within a time period t and feeding-in this energy, the feed-in amount of electricity should change the relevant RE account in *Layer 1*. The RE account enables participants to differentiate the types of electricity (renewable or not) in *Layer 2* especially regarding future market communication like the business processes to supply customers with RE levy free electricity or market processes for generating participants.

Actors: Three different actors are involved within this use case: households, providers, and the meter operator. Power generating prosumers need to reveal the trade-able amount of renewable energy when trading with the community. When billing consumer households, providers (utilities) need to identify the renewable electricity. The meter operator, who is in charge of the smart meter, has to provide the meter readings.

Prerequisites: Beside the general requirements regarding *Layer 0* mentioned above, we identified the following additional prerequisites:

1) The household has a smart meter that measures the accumulated feed-in to the external grid per time interval t.
2) The household has a smart meter that measures the amount of solar power generated per time interval t.
3) The measurement of the feed quantity for a time interval t is in t + 1 . Measurement results from t_0 to t_1 only exist from t_1.
4) For each smart meter, the meter operator provisions a measurement.

Trigger: The use case is triggered after the meter readings for a physical feed-in interval are transmitted. Thus, after the expiration of t_2.

Assumptions: Due to complex dependencies of external parameters, we made four assumptions:

1) Households cannot actively change meter readings and RE accounts after t_2.
2) Only renewable energy will be stored in the battery.
3) If not stated otherwise, households consume energy from energy sources with the priorities: local solar installation $\succ$ local battery $\succ$ external grid.

Procedure:

1) After expiration of t_1, the meter operator b_j creates a tuple for each smart meter of the feeding household h_i: (i) quantity $e_{i,t}$, in watt per hour (Wh) (ii) time interval, t_0 until t_1 and (iii) meter id m_i.
2) The measuring location sends the tuple to the meter operator. The operator provides an interface where a distributed component can collect the tuples for the Tokenization Contract within t_1 to t_2.
3) In t_2 till t_3, the Tokenization contract will update the RE account for the feed-in household. The calculation will be done according to a previously defined, with all participants attuned procedure.

Postconditions: An in-mutable record of the status of the household's RE account at time t exists in *Layer 1* that identifies an amount of electricity exempted from the RE levy.

Non-functional requirements: RE account information should not be tampered or deleted by any individual participant within the market, and should only be made available in a readable form to selected participants

Architectural Considerations: In the following, we retrospectively discuss individual architectural decisions concerning their impact on essential system qualities. For this use case, we had to decide on three essential architecture considerations: (I) Where should the blockchain node's execution environment be hosted? (II) Which software component commits smart meter data tuples to the blockchain? (III) Which actor should be held responsible for the blockchain nodes?

(I) Execution Environment of Blockchain Nodes: One option was to deploy the blockchain nodes onto available smart meters, but gauge appliance impedes software versions' delivery and makes them cost-intensive, and involves high maintenance efforts. On the other hand, available battery storage devices offered some limited software hosting capability. However, due to various interacting software components, the battery storage device's maintainability was evaluated as too expensive. Lastly, we decided to use a dedicated machine as a dedicated software component to minimize the degree of coupling for components. However, a comprehensive evaluation concerning a local or cloud-based deployment is still pending.

(II) Software Component for Meter Value Anchoring: Due to immutability requirements in this use case, consumption and generation values are measured and signed by the smart meter. While virtualizing full node blockchain clients on smart meter hardware is currently cost-intensive and involves high maintenance efforts due to the calibration process, we introduce a new software component called the Household Processing Unit (HPU). It provides the signed measured values for validation on the blockchain and later in the process sends the smart meter data to the netting entity. This minimizes the degree of coupling components and uses the HPU as a host for a full node blockchain client.

(III) Blockchain Node Authority: As the sole operator of all blockchain nodes, the smart meter operator is hardly suitable, as it already has control over all generation and consumption data. If this actor additionally commands all accounting data, liability issues could be a challenge. On the other hand, the smart meter operator is the only instance that can validate if households are committing correct meter values on the blockchain. Therefore, we propose distributing blockchain nodes fairly by deploying one in each household and at least one in the providers, and at least one within the smart meter operator IT infrastructure.

B. Accounting

Aim: Privacy-preserving on-chain data recording

Description: Trustworthy accounting of energy production and consumption in *Layer 0* requires relevant data to be provided and stored in a manipulation-safe way. This can be achieved by publishing smart meter data to *Layer 1* components. This establishes an immutable record of relevant data all interested parties can rely on for *Layer 2* services.

While this transparency may be desirable in some cases, information on energy consumption is highly sensitive and needs careful handling to meet privacy regulations and protect data providers. To address this problem, we provided in our earlier work [7] a proof- of-concept implementation using the ZoKrates framework [6] for verifiable off-chain computations and the Ethereum Blockchain. The basic idea is to provide hidden data as inputs to an off-chain program, which first computes a commitment to a smart meter value and assert that it equals the on-chain value, and then perform additional processing on the now validated data. After on-chain verification of this computation's correctness, the processing results can be trusted, even though the input data was never exposed.

Actors: Meter operator, household and utility are involved in this use case. The smart meter operator has to publish commitments to meter reading and serves data to authenticated households and utilities for validation.

Prerequisites: This use case relies on *Layer 0* elements for accessing smart meter data and *Layer 2* in order to compute verifiable off-chain computations, publish commitments to the blockchain and distribute data among participants.

Trigger: The publishing of commitments is triggered periodically. How data is served through interested parties is dependent on *Layer 2* implementation.

Assumptions: The *Layer 0* infrastructure is trustworthy.

Procedure:

1) On a trigger event, a commitment of data is published in *Layer 1*.
2) Commitments enable validation or further processing to authorized parties.
3) The authorized party validates the data by re-computing the commitment and comparing it to on-chain value. This step can be part of a verifiable off-chain program performing additional processing on the hidden data.

Postconditions: Authorized parties are convinced that the data they have been served is correct. Furthermore, they are convinced, that any processing result based on that data, which has been published to *Layer 1*, is also correct.

Architectural Considerations: This use case is primarily about privacy-preserving considerations of smart meter values on the blockchain. We reference this as an ”off-chain signatures pattern” [8]. Specifically, we do this via hash-anchoring. Therefore we deal with the following architectural questions: (I) How and where should the meter values be hashed? (II) How do off-chain components influence the agile development of a prototype? (III) How do we deal with delays or data losses for meter values during data transmission?

(I) Privacy-preserving Meter Values: To hash the meter values, we need a resource as in the tokenization use case. Due to the obstacle, as mentioned earlier of smart meters and the few and complexly interwoven battery storage devices, we decided to use dedicated hardware in the form of NUCs installed within the household environment (HPU). We are currently using sha256 as the hash method and calculating the hash with the following input values: timestamp, current meter value, and the household component's blockchain address. The algorithm's selection impacts the household component's performance and, thus, on the meter values' throughput. To achieve an improvement, we suggest optimizing the hashing function or using Pedersen commitments [9] in the future.

(II) Off-chaining and Agil Practices: We used agile methods in the prototype's development process and identified maintainability as an essential parameter in the architecture development. We noticed that some software components occasionally have to initialize artifacts during development, which is very time-consuming. This laborious process hinders our development flows, and the slicing of components becomes of great importance since subsequent changes are challenging to maintain.

(III) Smart Meter Value Caching: In order to enable privacy-preserving on-chain data recording, we have used off-chain components. For the computation carried out on these components, we periodically receive input values kept in memory. If the data is delayed or not received in an extensive network of participants, this leads to errors in the subsequent processing. Our solution for this challenge was to use a rolling cache. We have defined four-time slices of 15 minutes each to buffer one hour of missing or delayed values. This number proved to be a realistic parameter in our field test. Here, we want to refer to the trade-off not to define a too large cache. Otherwise, scaling problems could arise due to the meter value volume. However, it should not be too small either; otherwise, hash values cannot be calculated at all. Since the data is transmitted in clear text, we, furthermore, refer to the data minimization principle of the GDPR in article 5 (1(c)) [10] that states: ”Personal data shall be [...] limited to what is necessary in relation to the purposes for which they are processed”.

C. Netting

Aim: Transparent, law-compliant participation in RE levy exemption

Description: This use case describes the post-consumption activity at *Layer 1* of matching feed-ins at grid connection points (*Layer 0*) to withdrawers or assign each withdrawal at a network connection point (*Layer 0*) to one or more suppliers. When feeding in at a grid connection point, a working price, which is determined by the European Energy Exchange production price and additional electricity costs, is charged. If at least two participants in the community participate within a time period *t*, reduced ancillary electricity costs, as well as production prices, are possible due to a RE surcharge exemption if a proof of the community-generated solar energy is available.

Actors: Three different Actors are involved within this use case: households, providers, and the meter operator.

Prerequisites: Beside the general requirements regarding *Layer 0* and *Layer +1* mentioned above, we identified the following additional prerequisites:

1) The measurement of the withdrawal or feed-in quantity for a time interval t is available at t + 1 .Between time intervals t_1 and t_2, the measurement results from t_0 to t_1 are present.
2) For each smart meter, exactly one meter operator provides measurement via an interface.

Trigger: The Netting will be triggered automatically after the expiration of t_2.

Assumptions: Due to different distribution possibilities of energy quantities in *Layer 0*, we made three assumptions:

1) Households do not change actively their feed-in or withdrawal quantities.
2) We assume a post-consumption-billing. *Layer 2*
3) Although we assume static prices for energy distribution (end consumers always know prices beforehand), within *Layer 1* we can match dynamically supply and demand. Therefore we enable the underlying system of *Layer 2* to economically optimize the peer-to-peer trading.

Procedure:

1) After expiration of t_1 all smart meters sends their measurements regarding consumption and production of energy to an interface provided by the meter operator which collects all measurements for the Netting Contract.
2) In t_2 till t_3, the Netting contract will calculate the distribution of energy among the participants according to a previously defined, with all participants' attuned calculation method.
3) From the beginning of t_3, every household can retrieve their distribution results from the Netting contract which includes: (i) time interval (ii) smart meter id (iii) withdrawal or feed-in quantity, which was used within the community network (iv) withdrawal or feed-in quantity, which was not used within the community network.

Postconditions: All households can verify results from the Netting algorithm published to *Layer 1* and the peer-to-peer energy trading performed at *Layer 2* is economically optimized for the wealth of the community.

Architectural Considerations: As previously described [7], the netting entity is a dedicated component within our architecture responsible for calculating community-internal net production and consumption values. The chosen algorithm executes as a zero-knowledge verifiable off-chain computation and allows the verification on-chain without requiring access to sensitive meter values. Particularly for the netting entity, we concentrate on the following architectural questions: (I) What trade-offs exist between plenary and timely netting results? (II) How to handle the netting algorithm for our specific project requirements? (III) How to increase proof generation performance for a larger community scenario?

(I) Plenary vs. Timely Netting Results: One goal for this use case is to calculate netting results as timely as possible to immediately process the community's energy quantities. On the other hand, we need as much data as possible from individual participants to guarantee an optimized allocation of the community's energy data. However, individual households might, e.g., not participating in Netting due to network problems. This trade-off between completeness and timeliness represents a challenge for the netting calculation. Strict adherence to completeness would have effects on failure detection and handling. Too much generosity for calculation, however, also means a loss of time for prompt processing. Thus, we developed our algorithm so that all data must be transmitted to the Netting entity within 15 minutes and none-transmitting households do not participate in the respective billing period.

(II) Complex Netting Calculation: We used zero-knowledge proofs to guarantee trustworthiness. To that extend, we applied the ZoKrates Toolkit [6] to implement those proofs within our Ethereum development environment. However, to calculate Netting for creating virtual transfers between households, there are various approaches to allocate these energy flows. For example, we defined Consistency, Pareto Efficiency, and Fairness as our netting results' main properties. Due to scaling and specification challenges with ZoKrates, we calculate the Netting in a separate Python algorithm and extract the result for checking and validation using invariants for each property in ZoKrates. As a result, we create ZoKrates proofs, which are submitted on-chain and can be validated by interested parties.

(III) Serverless Netting: While adding households for testing our prototype, we experienced some performance challenges. The periodic execution of the ZoKrates program on a high-performance computer reaches our internal project computation limit of 15 minutes while calculating for 80 households. Therefore, we (I) switched from hash calculations with sha256 to Peterson Commits [9] and (II) implemented a distributed serverless execution model for scalable off-chain verifications by splitting assertions into sub-assertions for fast computation while maintaining completeness and light deployability.

D. Contracting

Aim: Predetermined matching of energy demand and supply

Description: This use case deals with the distribution of electricity before it is produced and consumed. Therefore, market participants (prosumer and consumer) automatically submit their offers and bids in a defined time window for a time interval announced in the future. The offers are composed of the forecasted values for consumption and generation of energy. We assume that several economic levers can be taken into account in Germany to realize a reduced electricity price for consumers within the community. For this purpose,

meta-information from *Layer 0* (e.g. location information of prosumer and consumer) must be added to the offers and bids. Due to this additional information, price-optimized pairings of producers and consumers by several cents per kilowatt-hour can be formed. The ex-ante consumption matching result is needed for creating a suggestion for the action of each household containing battery storage. This suggestion contains the information to either unload a certain amount of energy for the next given time window or optimize their self-consumption.
Actors: Two actors are involved in this use case: household and utility.
Prerequisites: For this use case, we composed three necessary prerequisites:

1) There is a forecast value for consumption and if applicable for the production of each household for a time interval announced in the future.
2) There is at least one utility offering power through the external grid in *Layer 0* for missing power supply bids.
3) There is at least one household that can be matched.

Assumptions: We made three assumptions:

1) Households with a battery storage, optimize their self-consumption.
2) We assume there is a service in *Layer 2* to get prediction values for the production and consumption of each household for every future time interval.
3) There exists a battery management interface where suggestions for action items can be transferred to.

Procedure:

1) Once triggered, forecast data concerning each participating household (production and consumption) will be received.
2) A calculation for the distribution of forecasted energy among the participants is done.
3) Forecasted data should be treated as described in *Accounting* (cf Sec. IV-B) due to personal information.
4) Suggestions for a battery management system should be only send when distribution results contain a energy providing household with a battery storage.

Postconditions: After the action statement has been sent to the battery management system, the matching data for the battery household should be stored privacy-conform on *Layer 1*, like in use case *Accounting*.

E. Compensation

Aim: Ensure correct compensation after violating energy trading contracts
Description: A producer (Alice) and a consumer (Bob) household can join in a energy trading contract c (cf Sec. IV-D). By doing so, Alice commits to feeding energy into the grid, and Bob commits to withdrawing energy from the grid, respectively. If one party violates c, the other party must use the utility as a fallback. However, trading with the utility results in a financial penalty. Thus, Alice and Bob deposit money in an escrow e managed by Charlie. The cost of using the utility determines the size of e. If Alice or Bob observes a violation of c, they can report it to Charlie. Charlie checks the violation based on smart meter readings. A similiar mechanism like the consent violation detection proposed in [11] can be used. If a violation occurs, Charlie changes the distribution of e.
Actors: This use case includes three actors. Alice commits to feeding energy into the public grid. Bob commits to withdrawing energy from the public grid. Charlie checks violations of energy contracts and releases escrowed funds.
Prerequisites: An energy contract c for period t, an escrow e.
Trigger: Alice or Bob triggers the use case by reporting a violation to Charlie.
Assumptions: The utility's prices are known (Layer +1). Thus, the size of e is deterministic. Alice and Bob can detect contract violations. Charlie can verify violations, and change the distribution of funds in e between Alice and Bob.
Procedure:

1) Alice observes Bob's violation of c (Layer 0).
2) Alice reports Bob's violation of c to Charlie (Layer 1).
3) Charlie verifies Alice's claim by checking Bob's commitment to c (Layer 1) and smart meter readings for t at Bob's public grid access point (Layer 0).
4) If the claim is verified, Charlie changes the distribution of e (Layer 1).
5) Charlie releases e (Layer 1).
6) Alice and Bob withdraw all funds from e (Layer 1).

Postconditions: Alice and Bob are financially compensated and e is empty.
Non-functional requirements: Charlies must provide processes in a correct, cost-efficient, scalable, and available manner (Layer +1).

V. RELATED WORK

Related surveys of Blockchain-based applications within the energy sector e.g., [1]–[4], and [12] show a great interest in implementing strong energy requirements with blockchain technologies. The interest is equally distributed among related work conducted by academia but also studies conducted by industry. Goranovic et al. [1] present a brief overview of 14 initiatives that develop applications for micro-grids using blockchain technology. Kim et al. [2] present a high-level comparison between private and public blockchains for electricity trading applications in micro-grids. Chitchyan and Murkin [3] give an introduction to 17 initiatives that develop blockchain-based applications for the energy sector. All three studies focus on industry-driven initiatives.

In contrast, Albrecht et al. [4] present a research method that includes expert interviews to assess the impact of blockchain technology on select use cases in the energy sector. Pipattanasomporn et al. [13] identify eight use cases embedded in peer-to-peer trading of solar energy and propose a Hyperledger-based prototype. The SINTEG project enera [14] published an experience report while developing more than 50 use cases.

Due to additional system objectives, recent publications focus on providing selected qualities, e.g., privacy [15], scalability [16], that constitute a non-trivial system design.

Other related proposals for energy trading enabled by blockchain technology range from simulations for Ethereum-based market place [13], [17], [18] to research has been done regarding battery storage within electric vehicle processes [19], [20]. But to the best of the authors' knowledge, none of them are presenting a detailed use case where combined battery storage could participate within an privacy-preserving energy trade.

VI. BENEFITS FROM BLOCKCHAIN

In the current energy market, actors, such as consumers, do not have to trust other actors. Participants only have to trust the calibration of their electricity meters. Their contracts rely on fixed prices for a long period of time. The introduced use case *Netting* (see section IV-C) is based on the idea of billing within 15 minute time windows via a utility with mixed prices and residual electricity for the anticipated community. This use case implies consumers to trust the utility. But since the utility incentive is to maximize the profit, it can't be trusted.

With our approach for a blockchain-based system, transparency and traceability can be established comparable to the current situation. This does not require a trust for a non-calibrated component and leads to a high acceptance of the consumer by mastering transparency and traceability. Also, it leads to the enablement of legally compliant audits. Another benefit is the elimination of a single point of failure. Furthermore, law compliant proof for external network participants, e.g., the main customs office (Hauptzollamt), can be created and made accessible without adding another trustworthy component.

VII. CONCLUSION

The current disruption within the energy market regarding the shift from a centralized, non-renewable to decentralized, renewable energy production poses many challenges. At the same time, it creates several opportunities, especially for private prosumers. In this paper, we have introduced a 3+1-Layer model that allows a structured development of trusted applications for this new kind of energy market. With the aid of this model, we achieved the challenging task of developing five relevant and sensible use cases and their implications on architectural decisions together with an industry consortium in the *BloGPV* project.

As an additional contribution besides the 3+1-Layer model, we have described and thoroughly discussed the identified use cases and have pointed out the architectural challenges regarding their realization. The analysis of those use cases has shown that concepts for their realization are missing, especially regarding the complex non-functional properties. We have presented how a blockchain-driven trusted computation layer can enable the realization of those use cases without a trusted third party. Still, some challenges, especially regarding privacy, are hard to fulfill, and new concepts to address them have to be developed. A promising approach is the use of zero-knowledge proofs to maintain the privacy that could be realized in a real-world system through the integration of the ZoKrates toolbox in *Layer 1*.

ACKNOWLEDGEMENT

The work in this paper was performed in the context of the BloGPV-Blossom project. *BloGPV-Blossom* is partially funded by the Germany Federal Ministry for Economic Affairs and Energy (BMWi) under grant no. 01MD18001E.

REFERENCES

[1] A. Goranović, M. Meisel, L. Fotiadis, S. Wilker, A. Treytl, and T. Sauter, "Blockchain applications in microgrids an overview of current projects and concepts," in *IECON 2017*, 2017, pp. 6153–6158.

[2] G. Kim, J. Park, and J. Ryou, "A study on utilization of blockchain for electricity trading in microgrid," in *BigComp'18*. IEEE, 2018.

[3] R. Chitchyan and J. Murkin, "Review of blockchain technology and its expectations: Case of the energy sector," *CoRR*, vol. abs/1803.03567, 2018. [Online]. Available: http://arxiv.org/abs/1803.03567

[4] S. Albrecht, S. Reichert, J. Schmid, J. Strüker, D. Neumann, and G. Fridgen, "Dynamics of blockchain implementation-a case study from the energy sector," in *Proc. of the 51st Hawaii Int. Conf. on System Sciences*, 2018.

[5] S. Tai, J. Eberhardt, and M. Klems, "Not ACID, not BASE, but SALT - a transaction processing perspective on blockchains," in *CLOSER'17*, INSTICC. ScitePress, 2017, pp. 755–764.

[6] J. Eberhardt and S. Tai, "Zokrates - scalable privacy-preserving off-chain computations," in *ICBC'18*. IEEE, 2018.

[7] J. Eberhardt, M. Peise, D.-H. Kim, and S. Tai, "Privacy-preserving netting in local energy grids," in *ICBC'20*. IEEE, 2020.

[8] J.Eberhardt and S.Tai, "On or Off the Blockchain? Insights on Off-Chaining Computation and Data," in *ESOCC 2017: 6th European Conference on Service-Oriented and Cloud Computing*, Oslo, Norway, Sep 2017. [Online]. Available: http://www.ise.tu-berlin.de/fileadmin/fg308/publications/2017/2017-eberhardt-tai-offchaining-patterns.pdf

[9] T. P. Pedersen, "Non-interactive and information-theoretic secure verifiable secret sharing," in *Advances in Cryptology — CRYPTO '91*, J. Feigenbaum, Ed. Berlin, Heidelberg: Springer Berlin Heidelberg, 1992, pp. 129–140.

[10] European Parliament & Council, "Regulation (EU) 2016/679 [...] on the protection of natural persons with regard to the processing of personal data and on the free movement of such data [...] (General Data Protection Regulation)," *Official Journal of the European Union*, vol. L 119/1, pp. 1–88, 2016.

[11] J. Heiss, M.-R. Ulbricht, and J. Eberhardt, "Put Your Money Where Your Mouth Is – Towards Blockchain-based Consent Violation Detection," in *ICBC 2020*. IEEE, 2020.

[12] M. Andoni, V. Robu, D. Flynn, S. Abram, D. Geach, D. Jenkins, P. McCallum, and A. Peacock, "Blockchain technology in the energy sector: A systematic review of challenges and opportunities," *Renewable and Sustainable Energy Reviews*, vol. 100, pp. 143 – 174, 2019.

[13] M. Pipattanasomporn, M. Kuzlu, and S. Rahman, "A blockchain-based platform for exchange of solar energy: Laboratory-scale implementation," in *ICUE'18*. IEEE, 2018.

[14] M. Clausen, R. Apel, M. Dorchain, M. Postina, and M. Uslar, "Use case methodology: a progress report," *Energy Informatics*, vol. 1, no. 1, 2018.

[15] F. Knirsch, A. Unterweger, and D. Engel, "Privacy-preserving blockchain-based electric vehicle charging with dynamic tariff decisions," vol. 33, no. 1, pp. 71–79, 2018.

[16] F. Blom and H. Farahmand, "On the scalability of blockchain-supported local energy markets," in *SEST'18*, 2018.

[17] C. Pop, T. Cioara, M. Antal, I. Anghel, I. Salomie, and M. Bertoncini, "Blockchain based decentralized management of demand response programs in smart energy grids," *Sensors*, vol. 18, no. 1, 2018.

[18] A. Hahn, R. Singh, C. Liu, and S. Chen, "Smart contract-based campus demonstration of decentralized transactive energy auctions," in *2017 IEEE Power Energy Society Innovative Smart Grid Technologies Conference (ISGT)*, 2017.

[19] X. Huang, C. Xu, P. Wang, and H. Liu, "Lnsc: A security model for electric vehicle and charging pile management based on blockchain ecosystem," *IEEE Access*, vol. 6, pp. 13 565–13 574, 2018.

[20] Z. Su, Y. Wang, Q. Xu, M. Fei, Y. Tian, and N. Zhang, "A secure charging scheme for electric vehicles with smart communities in energy blockchain," *IEEE Internet of Things Journal*, 2019.